

Neural Causal Abstractions

Kevin Xia and Elias Bareinboim

Causal Artificial Intelligence Lab
Columbia University
{kevinmxia, eb}@cs.columbia.edu

Abstract

The abilities of humans to understand the world in terms of cause and effect relationships, as well as to compress information into abstract concepts, are two hallmark features of human intelligence. These two topics have been studied in tandem in the literature under the rubric of causal abstractions theory. In practice, it remains an open problem how to best leverage abstraction theory in real-world causal inference tasks, where the true mechanisms are unknown and only limited data is available. In this paper, we develop a new family of causal abstractions by clustering variables and their domains. This approach refines and generalizes previous notions of abstractions to better accommodate individual causal distributions that are spawned by Pearl’s causal hierarchy. We show that such abstractions are learnable in practical settings through Neural Causal Models (Xia et al. 2021), enabling the use of the deep learning toolkit to solve various challenging causal inference tasks – identification, estimation, sampling – at different levels of granularity. Finally, we integrate these results with representation learning to create more flexible abstractions, moving these results closer to practical applications. Our experiments support the theory and illustrate how to scale causal inferences to high-dimensional settings involving image data.

1 Introduction

Humans understand the world around them through the use of abstract notions. Biologists can study the function of the liver without understanding the interactions between its subatomic particles studied by physicists. Economists find it more practical to consider macro-level behavior through concepts like aggregate supply and demand rather than studying the purchasing behavior of individuals. At home, we choose to interpret the object in the television as a dog or a car as opposed to a collection of photons or pixels. Humans are highly capable of learning through interacting with the environment and understanding cause and effect between different concepts. Understanding causality is considered a hallmark of human intelligence and allows humans to plan a course of action, determine blame and responsibility, and generalize across environments. It follows that the ability to abstract concepts and study them causally is a key ability expected from modern intelligent systems.

AI systems are built on a foundation of generative models, which are representations of the underlying processes

from which data is collected. Standard generative models simply model some joint density of a set of variables of interest, while *causal* generative models further model distributions involving causal interventions and counterfactual relations. In this paper, we study the problem of learning a causal generative model from data, which can be useful for many purposes such as sampling novel causally-consistent data points (i.e. from interventional or counterfactual distributions). One major challenge is that data is often provided in complex low level forms (e.g., pixels), while it would be more useful in applications to focus on higher level concepts (e.g., dog or car). We would therefore like to learn a more abstract causal generative model at a higher level of granularity, while guaranteeing that the queries from the coarser model match the ground truth.

To formalize this problem, we build on the semantics of a class of generative models called structural causal models (SCMs) (Pearl 2000). An SCM \mathcal{M}^* describes a collection of mechanisms and distribution over unobserved factors. Each SCM induces three qualitatively different sets of distributions related to the human concepts of “seeing” (called observational), “doing” (interventional), and “imagining” (counterfactual), collectively known as the Ladder of Causation or the Pearl Causal Hierarchy (PCH) (Pearl and Mackenzie 2018; Bareinboim et al. 2022). The PCH is a containment hierarchy in which each of these distribution sets can be put into increasingly refined layers, where observational distributions go in layer 1 (\mathcal{L}_1), interventional in layer 2 (\mathcal{L}_2), and counterfactual in layer 3 (\mathcal{L}_3). In typical tasks of causal inference, the goal is to obtain a quantity from a higher layer when given data only from lower layers (e.g. inferring interventional quantities from observational data). Still, it is understood that this is generally impossible without additional assumptions since higher layers are underdetermined by lower layers (Bareinboim et al. 2022; Ibeling and Icard 2020).

Generative models can often be implemented in practice as neural networks. Deep learning models have achieved promising success in a variety of applications such as computer vision (Krizhevsky, Sutskever, and Hinton 2012), speech recognition (Graves and Jaitly 2014), and game playing (Mnih et al. 2013). Many of these successes are attributed to *representation learning* (Bengio, Courville, and Vincent 2013), in which the learned representation can be thought of as an abstraction of the data. Further, there has also been grow-

ing interest in the idea of incorporating causality into deep models¹. Prior work introduced one such model, the Neural Causal Model (NCM), which incorporates the same causal assumptions encoded in a causal diagram to identify and estimate interventional and counterfactual distributions (Xia et al. 2021; Xia, Pan, and Bareinboim 2023).² Despite the soundness of this approach in theory, current NCM-based methods face challenges when applied to complex real-world settings for various reasons: (1) optimization is difficult when scaled to high dimensions, (2) unprocessed data can come in complicated forms (e.g. images, text, etc.), and (3) the causal diagram is difficult to fully specify in some high-dimensional settings. In this paper, we address these challenges by studying how representation learning and causal reasoning are related to each other and by building on this understanding to develop a neural framework for causal abstraction learning.

Existing works that study causal abstractions set a solid foundation by defining various mathematical notions of abstractions (Rubenstein et al. 2017; Beckers and Halpern 2019; Beckers, Eberhardt, and Halpern 2019; Geiger, Potts, and Icard 2023; Massidda et al. 2023). In App. B, we explain some of the foundational results and discuss their drawbacks. In particular, we note that existing definitions are declarative; that is, if the lower and higher level models are given, one can use the definition to decide whether the higher level model is indeed an abstraction of the lower level one. However, neither models are available in practice, and one would want to use limited lower level data to learn a higher level causal abstraction. We will expand on the current generation of causal abstractions in two ways. First, given that the true SCM is almost never available in practice, nor entirely learnable from data, we introduce a relaxed notion of abstractions that applies on the layers of the PCH. Second, we develop algorithms to systematically learn abstractions in practice given some structural information about the data, which can then be used for downstream inferential tasks such as causal identification, estimation, or sampling.

The general problem tackled by this paper is summarized in Fig. 1. The ground truth model \mathcal{M}_L (left) is defined over low level variables \mathbf{V}_L (e.g., pixels), while it may be practical to work in their high level abstract counterparts \mathbf{V}_H (e.g., dog or car). \mathcal{M}_L induces distributions from the three layers of the PCH (i.e. \mathcal{L}_1^* , \mathcal{L}_2^* , \mathcal{L}_3^*), defined over \mathbf{V}_L . In this work, we introduce a new type of abstraction function τ that maps distributions over \mathbf{V}_L to ones over \mathbf{V}_H (i.e. $\tau(\mathcal{L}_1^*)$, $\tau(\mathcal{L}_2^*)$, $\tau(\mathcal{L}_3^*)$). Furthermore, \mathcal{M}_L is unobserved, and only limited data is given (e.g., observational data from

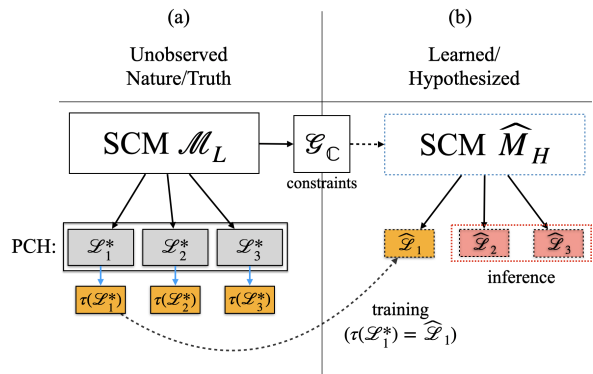


Figure 1: Overview of this paper. High-level SCM $\widehat{\mathcal{M}}_H$ (right) is trained on available data to serve as an abstract proxy of the true, unobserved, low-level SCM \mathcal{M}_L (left).

\mathcal{L}_1^*). The goal is to learn a high-level SCM $\widehat{\mathcal{M}}_H$ (right) over the high-level variables \mathbf{V}_H that encodes the given causal constraints (\mathcal{G}_C in the figure) and matches \mathcal{M}_L on the available data across τ (e.g. $\widehat{\mathcal{L}}_1 = \tau(\mathcal{L}_1^*)$). Then, we investigate when and how the resulting model $\widehat{\mathcal{M}}_H$ can be used as a surrogate, allowing one to make interventional and counterfactual inferences about the higher layers of \mathcal{M}_L through the higher layers of $\widehat{\mathcal{M}}_H$.

As an example, suppose an economist is studying the effects of spending trends of various countries on their average income Y . In addition to Y , she has collected observational data on several variables of spending trends, such as consumer spending C , investments I , government spending G , imports M , and exports E . She wants to understand the causal effect of increasing general spending on average income of the population, and one way to do this is to study the causal effect of collectively increasing C , I , G , M , and E on Y . However, the economist notes that C , I , G , M , and E can be aggregated together into a single abstract variable called gross domestic product (GDP). The tools that we introduce in this paper allow her to proceed by constructing a high-level model $\widehat{\mathcal{M}}_H$ over the variables GDP and Y , encoding the required causal assumptions, and training the model over the given observational data. Despite the high-level variables not matching the original low-level variables (i.e., C , I , G , M , E), the causal effect of GDP on Y can be queried from the model $\widehat{\mathcal{M}}_H$ to solve the problem.

More specifically, our contributions are as follows: In Sec. 2, we define a new class of abstractions based on clusters of variables (intervariable) and their domains (intravariable). Building on this new class, we define a notion of abstraction consistency on the layers of the PCH. We then show how to systematically construct an abstraction consistent with all three layers of the PCH and then relate these abstractions to existing definitions. In Sec. 3, we show how to leverage NCM machinery to perform interventional (layer 2) and counterfactual (layer 3) inferences across these abstractions when the true SCM is unavailable. In Sec. 4, we introduce a variant of the NCM that learns representations of each variable

¹Many successful approaches have been developed to estimate causal effects from observational data under the backdoor or conditional ignorability conditions (Shalit, Johansson, and Sontag 2017; Louizos et al. 2017; Li and Fu 2017; Johansson, Shalit, and Sontag 2016; Yao et al. 2018; Yoon, Jordon, and van der Schaar 2018; Kallus 2020; Shi, Blei, and Veitch 2019; Du et al. 2020; Guo et al. 2020), and also to answer causal queries through neural-parameterized SCMs (Kocaoglu et al. 2018; Goudet et al. 2018).

²The literature also includes non-neural approaches for such problems, including estimators with stronger statistical properties such as double robustness and convergence guarantees, for example, (Jung, Tian, and Bareinboim 2020a,b, 2021).

and encodes causal assumptions on the representation level, allowing us to learn abstractions even in settings where the assumption of the availability of clusters is relaxed. Experiments in Sec. 5 corroborate with the theory.

1.1 Preliminaries

We now introduce the notation and definitions used throughout the paper. We use uppercase letters (X) to denote random variables and lowercase letters (x) to denote corresponding values. Similarly, bold uppercase (\mathbf{X}) and lowercase (\mathbf{x}) letters denote sets of random variables and values respectively. We use \mathcal{D}_X to denote the domain of X and $\mathcal{D}_{\mathbf{X}} = \mathcal{D}_{X_1} \times \dots \times \mathcal{D}_{X_k}$ for the domain of $\mathbf{X} = \{X_1, \dots, X_k\}$. We denote $P(\mathbf{X} = \mathbf{x})$ (often shortened to $P(\mathbf{x})$) as the probability of \mathbf{X} taking the values \mathbf{x} under the distribution $P(\mathbf{X})$.

We utilize the basic semantic framework of structural causal models (SCMs) (Pearl 2000), following the presentation in Bareinboim et al. (2022).

Definition 1 (Structural Causal Model (SCM)). A structural causal model \mathcal{M} is a 4-tuple $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$, where

- \mathbf{U} is a set of background (exogenous) variables that are determined by factors outside the model;
- \mathbf{V} is a set $\{V_1, V_2, \dots, V_n\}$ of variables, called endogenous, that are determined by other variables in the model – that is, variables in $\mathbf{U} \cup \mathbf{V}$;
- \mathcal{F} is a set of functions $\{f_{V_1}, f_{V_2}, \dots, f_{V_n}\}$ such that each f_{V_i} is a mapping from exogenous parents $\mathbf{U}_{V_i} \subseteq \mathbf{U}$ and endogenous parents $\mathbf{Pa}_{V_i} \subseteq \mathbf{V} \setminus V_i$ to V_i ;
- $P(\mathbf{U})$ is a probability function defined over $\mathcal{D}_{\mathbf{U}}$. ■

Definition 2 (Causal Diagram (Bareinboim et al. 2022, Def. 13)). Each SCM \mathcal{M} induces a causal diagram \mathcal{G} , constructed as follows:

1. add a vertex for each $V_i \in \mathbf{V}$;
2. add a directed arrow ($V_j \rightarrow V_i$) for every $V_i \in \mathbf{V}$ and $V_j \in \mathbf{Pa}_{V_i}$; and
3. add a dashed-bidirected arrow ($V_j \dashleftrightarrow V_i$) for every pair $V_i, V_j \in \mathbf{V}$ such that \mathbf{U}_{V_i} and \mathbf{U}_{V_j} are not independent (Markovianity is not assumed). ■

Our treatment is constrained to *recursive* SCMs, which implies acyclic causal diagrams, with finite discrete domains over endogenous variables \mathbf{V} .

Counterfactual (and also interventional and observational) quantities can be computed from SCM \mathcal{M} as follows:

Definition 3 (Layer 3 Valuation (Bareinboim et al. 2022, Def. 7)). An SCM \mathcal{M} induces layer $\mathcal{L}_3(\mathcal{M})$, a set of distributions over \mathbf{V} , each with the form $P(\mathbf{Y}_*) = P(\mathbf{Y}_{1[x_1]}, \mathbf{Y}_{2[x_2]}, \dots)$ such that

$$P^{\mathcal{M}}(\mathbf{y}_{1[x_1]}, \mathbf{y}_{2[x_2]}, \dots) = \int_{\mathcal{D}_{\mathbf{U}}} \mathbf{1}[\mathbf{Y}_{1[x_1]}(\mathbf{u}) = \mathbf{y}_1, \mathbf{Y}_{2[x_2]}(\mathbf{u}) = \mathbf{y}_2, \dots] dP(\mathbf{u}) \quad (1)$$

where $\mathbf{Y}_{i[x_i]}(\mathbf{u})$ is evaluated under $\mathcal{F}_{\mathbf{x}_i} := \{f_{V_j} : V_j \in \mathbf{V} \setminus \mathbf{X}_i\} \cup \{f_X \leftarrow x : X \in \mathbf{X}_i\}$. \mathcal{L}_2 is the subset of \mathcal{L}_3 for which all \mathbf{x}_i are equal, and \mathcal{L}_1 is the subset for which all $\mathbf{X}_i = \emptyset$. ■

Each \mathbf{Y}_i corresponds to a set of variables in a world where the original mechanisms f_X are replaced with constants \mathbf{x}_i for each $X \in \mathbf{X}_i$; this is also known as the mutilation procedure. This procedure corresponds to interventions, and we use subscripts to denote the intervening variables (e.g. $\mathbf{Y}_{\mathbf{x}}$) or subscripts with brackets when the variables are indexed (e.g. $\mathbf{Y}_{1[x_1]}$). For instance, $P(y_x, y_{x'})$ is the probability of the joint counterfactual event $Y = y$ had X been x and $Y = y'$ had X been x' .

We use the notation $\mathcal{L}_i(\mathcal{M})$ to denote the set of \mathcal{L}_i distributions from \mathcal{M} . We use \mathbb{Z} to denote a set of quantities from Layer 2 (i.e. $\mathbb{Z} = \{P(\mathbf{V}_{\mathbf{z}_k})\}_{k=1}^{\ell}$), and $\mathbb{Z}(\mathcal{M})$ denotes those same quantities induced by SCM \mathcal{M} (i.e. $\mathbb{Z}(\mathcal{M}) = \{P^{\mathcal{M}}(\mathbf{V}_{\mathbf{z}_k})\}_{k=1}^{\ell}$).

We also build on Neural Causal Models (NCMs), in particular for performing causal inferences:

Definition 4 (\mathcal{G} -Constrained Neural Causal Model (\mathcal{G} -NCM) (Xia et al. 2021, Def. 7)). Given a causal diagram \mathcal{G} , a \mathcal{G} -constrained Neural Causal Model (for short, \mathcal{G} -NCM) $\widehat{\mathcal{M}}(\theta)$ over variables \mathbf{V} with parameters $\theta = \{\theta_{V_i} : V_i \in \mathbf{V}\}$ is an SCM $\langle \widehat{\mathbf{U}}, \mathbf{V}, \widehat{\mathcal{F}}, P(\widehat{\mathbf{U}}) \rangle$ such that

- $\widehat{\mathbf{U}} = \{\widehat{U}_{\mathbf{C}} : \mathbf{C} \in \mathbb{C}(\mathcal{G})\}$, where $\mathbb{C}(\mathcal{G})$ is the set of all maximal cliques over bidirected edges of \mathcal{G} ;
- $\widehat{\mathcal{F}} = \{\widehat{f}_{V_i} : V_i \in \mathbf{V}\}$, where each \widehat{f}_{V_i} is a feedforward neural network parameterized by $\theta_{V_i} \in \theta$ mapping values of $\mathbf{U}_{V_i} \cup \mathbf{Pa}_{V_i}$ to values of V_i for $\mathbf{U}_{V_i} = \{\widehat{U}_{\mathbf{C}} : \widehat{U}_{\mathbf{C}} \in \widehat{\mathbf{U}} \text{ s.t. } V_i \in \mathbf{C}\}$ and $\mathbf{Pa}_{V_i} = \text{Pa}_{\mathcal{G}}(V_i)$;
- $P(\widehat{\mathbf{U}})$ is defined s.t. $\widehat{U} \sim \text{Unif}(0, 1)$ for each $\widehat{U} \in \widehat{\mathbf{U}}$. ■

2 Abstractions of the Pearl Causal Hierarchy

The discussion of abstractions begins with defining causal variables. In many established causal inference tasks, it is typically assumed that there is a well-specified and known set of endogenous variables of interest \mathbf{V} , and nature is modeled by a collection of mechanisms that assign values to each of these variables. However, in practice, the definition of \mathbf{V} may not always be clear. In particular, the variables of interest may not align with the features of the data. For example, in an economic system, perhaps data on each individual consumer is collected, but the variable of interest is an aggregate measure like gross domestic product (GDP). In image data, perhaps the pixel values are collected, but the variables of interest are related to the objects of the image, not the individual pixels.

Acknowledging that the data is not always provided in the best choice of granularity, the causal abstraction literature typically defines two sets of variables, \mathbf{V}_L and \mathbf{V}_H , which describe the lower level and higher level settings, respectively. For example, \mathbf{V}_L might describe the pixels of an image, while \mathbf{V}_H might describe its structural content. They are typically modeled by corresponding causal models \mathcal{M}_L and \mathcal{M}_H , respectively.

In this section, we study on the distinction between low level variables \mathbf{V}_L (e.g. pixels) and their higher level counterparts \mathbf{V}_H (e.g. image) from the perspective of individual distributions of the PCH. We consider nature’s underlying SCM \mathcal{M}_L defined over low level variables, \mathbf{V}_L , and the goal

is to reason about the higher level variables \mathbf{V}_H given data on \mathbf{V}_L ³.

2.1 Constructive Abstraction Functions

The connection between \mathbf{V}_H and \mathbf{V}_L can be described through a mapping between their domains, $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$. Here, we consider a family of abstraction functions where τ is based on clusters of the variables and values of \mathbf{V}_L :

Definition 5 (Inter/Intravariabe Clustering). Let \mathcal{M} be an SCM over variables \mathbf{V} .

1. A set \mathbb{C} is said to be an intervariable clustering of \mathbf{V} if $\mathbb{C} = \{\mathbb{C}_1, \mathbb{C}_2, \dots, \mathbb{C}_n\}$ is a partition of a subset of \mathbf{V} . \mathbb{C} is further considered admissible w.r.t. \mathcal{M} if for any $\mathbb{C}_i \in \mathbb{C}$ and any $V \in \mathbb{C}_i$, no descendent of V outside of \mathbb{C}_i is an ancestor of any variable in \mathbb{C}_i . That is, there exists a topological ordering of the clusters of \mathbb{C} relative to the functions of \mathcal{M} .
2. A set \mathbb{D} is said to be an intravariabe clustering of variables \mathbf{V} w.r.t. \mathbb{C} if $\mathbb{D} = \{\mathbb{D}_{\mathbb{C}_i} : \mathbb{C}_i \in \mathbb{C}\}$, where $\mathbb{D}_{\mathbb{C}_i} = \{\mathcal{D}_{\mathbb{C}_i}^1, \mathcal{D}_{\mathbb{C}_i}^2, \dots, \mathcal{D}_{\mathbb{C}_i}^{m_i}\}$ is a partition (of size m_i) of the domains of the variables in \mathbb{C}_i , $\mathcal{D}_{\mathbb{C}_i}$ (recall that $\mathcal{D}_{\mathbb{C}_i}$ is the Cartesian product $\mathcal{D}_{V_1} \times \mathcal{D}_{V_2} \times \dots \times \mathcal{D}_{V_k}$ for $\mathbb{C}_i = \{V_1, V_2, \dots, V_k\}$, so elements of $\mathcal{D}_{\mathbb{C}_i}^j$ take the form of tuples of the value settings of \mathbb{C}_i). ■

In words, \mathbf{V} is divided into n subsets or clusters $\mathbb{C}_1, \dots, \mathbb{C}_n$ (variables that are not put into one of the clusters are projected away), and they are called *intervariable* clusters because the variables themselves are divided apart. Admissibility implies that the recursivity assumption of SCMs is retained through the intervariable clusters. Then, the joint *domains* of each of these n clusters are further partitioned. For example, for a specific intervariable cluster \mathbb{C}_i , the domain $\mathcal{D}_{\mathbb{C}_i}$ contains the set of all tuples of values of \mathbb{C}_i , and $\mathbb{D}_{\mathbb{C}_i}$ describes a partition $\mathcal{D}_{\mathbb{C}_i}^1, \dots, \mathcal{D}_{\mathbb{C}_i}^{m_i}$ of size m_i over this set of values (i.e. each $\mathcal{D}_{\mathbb{C}_i}^j \subseteq \mathcal{D}_{\mathbb{C}_i}$). The *intravariabe* clusters are the set of the value partitions over each intervariable cluster, and the term ‘‘intravariabe’’ denotes that the clustering is within the variable domains. Intuitively speaking, intervariable clusters partition the low level variables to describe each high level variable as a collection of low level variables. Intravariabe clusters then describe the domains of these high level variables by partitioning the corresponding value spaces of these intervariable clusters.

Example 1. Consider a study on the effects of certain food dishes on body mass index (BMI), inspired by nutrition studies like Gamba et al. (2014). Data is collected on individuals eating at restaurants, including the restaurant (R), dish ordered (D), the amount of carbohydrates (C), fat (F), and protein (P) in the dish, and the BMI of the customer (B). That is, $\mathbf{V}_L = \{R, D, C, F, P, B\}$. One food scientist argues that any nutritional impact of the food on BMI could be abstracted based on how many calories are in each dish. One may then be tempted to cluster the variables C , F , and P

³For concreteness, we assume that \mathcal{M}_L is an SCM, but the underlying generative model can be left implicit as explained in Appendix D.1.




Figure 2: Example of a constructive abstraction function τ w.r.t. corresponding inter/intravariabe clusters. **Top (inter-variable):** The low-level variables, dish (D) and BMI (B), are in their own clusters while restaurant (R) is abstracted away. Carbohydrates (C), fat (F), and protein (P) are clustered together and are mapped to a single variable, calories (Z). **Bottom (intravariabe):** The intravariabe clustering for $\mathbb{C}_2 = \{C, F, P\}$ is shown. Calories Z can be computed from C, F, P using the formula $Z = 4C + 9F + 4P$. This means that the domain is partitioned such that two different values, $(c_1, f_1, p_1), (c_2, f_2, p_2)$ are in the same intravariabe cluster if $4c_1 + 9f_1 + 4p_1 = 4c_2 + 9f_2 + 4p_2$.

together into one variable, named calories, labeled Z . This is an example of intervariable clustering.

To denote this formally, we may choose $\mathbb{C} = \{\mathbb{C}_1 = \{B\}, \mathbb{C}_2 = \{C, F, P\}, \mathbb{C}_3 = \{D\}\}$ as the intervariable clusters. In this case, B and D are placed in their own clusters, \mathbb{C}_1 and \mathbb{C}_3 , respectively. C, F , and P are all clustered together into \mathbb{C}_2 . R is not included and is abstracted away, which may be desirable if R is not relevant to the study. Collectively, $\mathbb{C}_1, \mathbb{C}_2$, and \mathbb{C}_3 form a partition of the subset of \mathbf{V}_L without R . Each of the clusters of \mathbb{C} will correspond to a high level variable of \mathbf{V}_H . In this case, for example, let Z denote the high level variable corresponding to cluster \mathbb{C}_2 , interpreted as calories. This is shown at the top of Fig. 2 (red).

The domain of \mathbb{C}_2 contains every tuple of C, F , and P , but the domain of Z can be simplified. After all, the computation of calories can be specified as $Z = 4C + 9F + 4P$, which means that two sets of values, $(c_1, f_1, p_1), (c_2, f_2, p_2)$ are considered equivalent if $4c_1 + 9f_1 + 4p_1 = 4c_2 + 9f_2 + 4p_2$. This clustering of domain values is an example of intravariabe clustering, shown at the bottom of Fig. 2 (blue). More formally, the intervariable clusters would be denoted $\mathbb{D} = \{\mathbb{D}_{\mathbb{C}_1}, \mathbb{D}_{\mathbb{C}_2}, \mathbb{D}_{\mathbb{C}_3}\}$, where each $\mathbb{D}_{\mathbb{C}_i}$ is a partition of $\mathcal{D}_{\mathbb{C}_i}$. In the case of $\mathbb{D}_{\mathbb{C}_2}$, we may define $\mathbb{D}_{\mathbb{C}_2} = \{\mathcal{D}_{\mathbb{C}_2}^1, \mathcal{D}_{\mathbb{C}_2}^2, \dots\}$, where each $\mathcal{D}_{\mathbb{C}_2}^j$ is a collection of tuples $(c, f, p) \in \mathcal{D}_{\mathbb{C}_2}$ corresponding to some specific value $4c + 9f + 4p$. In Fig. 2 for example, $\mathcal{D}_{\mathbb{C}_2}^1 = \{(c, f, p) : 4c + 9f + 4p = 200, (c, f, p) \in \mathcal{D}_{\mathbb{C}_2}\}$. Each of the intravariabe clusters correspond to a domain value of the high level variable. For example, $\mathcal{D}_{\mathbb{C}_2}^1$ corresponds to a value of $Z = 200$. ■

For the remainder of this paper, we consider settings where the intervariable clusters are admissible. Collectively, given an intervariable clustering \mathbb{C} and intravariabile clustering \mathbb{D} of \mathbf{V}_L , an abstraction function τ can be defined as follows.

Definition 6 (Constructive Abstraction Function). A function $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ is said to be a constructive abstraction function w.r.t. inter/intravariabile clusters \mathbb{C} and \mathbb{D} iff

1. There exists a bijective mapping between \mathbf{V}_H and \mathbb{C} such that each $V_{H,i} \in \mathbf{V}_H$ corresponds to $\mathbf{C}_i \in \mathbb{C}$;
2. For each $V_{H,i} \in \mathbf{V}_H$, there exists a bijective mapping between $\mathcal{D}_{V_{H,i}}$ and $\mathbb{D}_{\mathbf{C}_i}$ such that each $v_{H,i}^j \in \mathcal{D}_{V_{H,i}}$ corresponds to $\mathcal{D}_{\mathbf{C}_i}^j \in \mathbb{D}_{\mathbf{C}_i}$; and
3. τ is composed of subfunctions $\tau_{\mathbf{C}_i}$ for each $\mathbf{C}_i \in \mathbb{C}$ such that $\mathbf{v}_H = \tau(\mathbf{v}_L) = (\tau_{\mathbf{C}_i}(\mathbf{c}_i) : \mathbf{C}_i \in \mathbb{C})$, where $\tau_{\mathbf{C}_i}(\mathbf{c}_i) = v_{H,i}^j$ if and only if $\mathbf{c}_i \in \mathcal{D}_{\mathbf{C}_i}^j$. We also apply the same notation for any $\mathbf{W}_L \subseteq \mathbf{V}_L$ such that \mathbf{W}_L is a union of clusters in \mathbb{C} (i.e. $\tau(\mathbf{w}_L) = (\tau_{\mathbf{C}_i}(\mathbf{c}_i) : \mathbf{C}_i \in \mathbb{C}, \mathbf{C}_i \subseteq \mathbf{W}_L)$). ■

In words, through the subfunction $\tau_{\mathbf{C}_i}$, each low level cluster $\mathbf{C}_i \in \mathbb{C}$ maps to a single high level variable $V_{H,i} \in \mathbf{V}_H$, and the value $\mathbf{c}_i \in \mathcal{D}_{\mathbf{C}_i}$ maps to a corresponding high level value $v_{H,i}^j \in \mathcal{D}_{V_{H,i}}$. Specifically, $\tau_{\mathbf{C}_i}(\mathbf{c}_i)$ maps to $v_{H,i}^j$ if \mathbf{c}_i is in the intravariabile cluster $\mathcal{D}_{\mathbf{C}_i}^j$. Then, the overall function τ is simply composed of the subfunctions $\tau_{\mathbf{C}_i}$. Intuitively, τ is a constructive abstraction function if it maps \mathbf{V}_L to \mathbf{V}_H by first grouping the variables w.r.t. their corresponding intervariable cluster in \mathbb{C} (red maps to yellow in Fig. 2 (top)), followed by assigning each cluster a value based on which intravariabile cluster they belong in \mathbb{D} (blue maps to green in Fig. 2 (bottom)). As a result, \mathbf{V}_H can be interpreted such that $\mathbf{V}_H = \mathbb{C}$ and $\mathcal{D}_{V_{H,i}} = \mathbb{D}_{\mathbf{C}_i}$ for each $V_{H,i} \in \mathbf{V}_H$ ⁴. This construction of τ means that τ is unique given the clusters \mathbb{C} and \mathbb{D} (up to a renaming of the variables \mathbf{V}_H and its values $\mathcal{D}_{\mathbf{V}_H}$).

Example 2 (Example 1 continued). Suppose the high level variables are denoted as $\mathbf{V}_H = \{D_H, Z, B_H\}$, where D_H and B_H represent the high level counterparts of D and B that remain unchanged across the abstraction. Each high-level variable (i.e. D_H, Z, B_H) corresponds to an intervariable cluster (i.e., $\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3$, respectively), establishing a bijective connection between \mathbf{V}_H and \mathbb{C} . Each of their domains also correspond to an intravariabile cluster in \mathbb{D} . For example, each value of $Z = z$ corresponds to the choice of $\mathcal{D}_{\mathbf{C}_2}^j$ such that $4c + 9f + 4p = z$. The constructive abstraction function τ constructed from the clusters \mathbb{C} and \mathbb{D} would map $(D_H, Z, B_H) \leftarrow \tau(R, D, C, F, P, D)$, which can be decomposed as

$$\begin{aligned} \tau(R, D, C, F, P, D) &= (\tau_{\mathbf{C}_1}(D), \tau_{\mathbf{C}_2}(C, F, P), \tau_{\mathbf{C}_3}(B)) & (2) \\ &= (D, 4C + 9F + 4P, B). & (3) \end{aligned}$$

⁴For another example of abstractions constructed from clusters, see App. E Ex. 18. For examples of abstraction functions that are not constructive, see Ex. 20.

Observe that τ is broken down into $\tau_{\mathbf{C}_1}$, $\tau_{\mathbf{C}_2}$, and $\tau_{\mathbf{C}_3}$, which maps the variables of each intervariable cluster to their corresponding high level variable D_H, Z , and B_H , respectively. D and B are not affected by the abstraction in this example, so $\tau_{\mathbf{C}_1}$ and $\tau_{\mathbf{C}_3}$ are the identity function, directly setting $D_H \leftarrow D$ and $B_H \leftarrow B$. However, the calories, Z , is computed through $\tau_{\mathbf{C}_2}(C, F, P) = 4C + 9F + 4P$. This ensures that all values from $\mathcal{D}_{\mathbf{C}_2} = \mathcal{D}_C \times \mathcal{D}_F \times \mathcal{D}_P$ that are in the same intravariabile cluster are mapped to the same value of Z . This mapping is illustrated in Fig. 2. ■

Note that the relationship between \mathbf{V}_L and \mathbf{V}_H modeled by τ is not causal. Rather, the contents of \mathbf{V}_L constitute \mathbf{V}_H ⁵. Intuitively, two variables of \mathbf{V}_L are mapped to the same intervariable cluster if they constitute the same high level variable (e.g. two pixels of the same dog), and two values are mapped to the same intravariabile cluster if, from a higher level perspective, they are functionally identical (e.g. same image of the dog but rotated or cropped). In this sense, intravariabile clustering can be thought of as invariances in the data, since downstream functions are *invariant* to values that are in the same intravariabile cluster⁶.

This paper will focus on abstractions based on constructive abstraction functions τ created from intervariable and intravariabile clusters. This is in contrast with the previous works on causal abstractions discussed in App. B, which leave the functional form of τ implicit.⁷ One benefit of making τ concrete is that it allows for a rigorous definition of equivalence between the distributions of a low level model and that of a high level model, as will be elaborated next.

2.2 Layer-Specific Abstractions

Ultimately, we would like to study causal properties of \mathbf{V}_L through their higher level counterparts \mathbf{V}_H . A sensible goal is, therefore, to learn an SCM \mathcal{M}_H over \mathbf{V}_H , which can then be queried for causal inference tasks. Still, even if \mathbf{V}_H and \mathbf{V}_L are connected through some function τ , this alone does not imply that \mathcal{M}_H is an abstraction of \mathcal{M}_L . This is the case since the distributions over \mathbf{V}_H induced by \mathcal{M}_H may not have any clear connection with the distributions over \mathbf{V}_L . To explain this point with a simple example, suppose $\mathbf{V}_L = \{X, Y\}$ and $\mathbf{V}_H = \mathbf{V}_L$ (i.e., τ is the identity function). However, in \mathcal{M}_L , f_X takes Y as an argument, while in \mathcal{M}_H , f_Y takes X as an argument. These two models are obviously unrelated despite sharing the same set of variables.

When two SCMs are defined over the same space of variables, one can verify that they are similar if they induce the same distributions. For example, an SCM \mathcal{M}' is \mathcal{L}_2 -consistent with \mathcal{M} if $\mathcal{L}_2(\mathcal{M}') = \mathcal{L}_2(\mathcal{M})$, that is, \mathcal{M} and \mathcal{M}' match in every interventional distribution (Bareinboim et al. 2022; Xia et al. 2021). However, when two SCMs are

⁵The distinction between causal and constitutional relationships is important and is explained in detail in Appendix D.1.

⁶This analogy is explored further in Apps. D.2 and D.3.

⁷In a related setting, there may be cases where information about \mathcal{M}_H is provided while τ is unknown. In such cases, there are works in the literature that aim to learn τ (Zennaro et al. 2023; Felekis et al. 2024) such that desired abstraction properties hold. Such functions may or may not be constructive, refer to App. B.2 for further details.

defined over different variable spaces, comparing their distributions is no longer well-defined. Hence, a different notion of consistency is needed to compare an SCM over \mathbf{V}_L with another over \mathbf{V}_H through τ .

We first note that not all low-level quantities have a clear interpretation in the high-level setting. For instance, in Example 1, the low-level quantity $P(C = 5)$ does not have a counterpart in \mathbf{V}_H because C is clustered together with F and P in the intervariable clusters. As another example, $P(C = 5, F = 0, P = 0)$ also does not have a counterpart in \mathbf{V}_H because τ maps $(C = 5, F = 0, P = 0)$ to the same value of Z as other quantities such as $(C = 0, F = 0, P = 5)$. In both cases, $Z = 4C + 9F + 4P = 20$. Hence, $P(C = 5, F = 0, P = 0)$ is not equivalent to $P(Z = 20)$, which, in a sense, represents all values of (C, F, P) such that $4C + 9F + 4P = 20$.

To define the low-level counterfactual quantities that have high-level counterparts through τ , first denote $\mathbf{Y}_{L,*}$ as a set of counterfactual variables over \mathbf{V}_L . That is,

$$\mathbf{Y}_{L,*} = (\mathbf{Y}_{L,1[\mathbf{x}_{L,1}]}, \mathbf{Y}_{L,2[\mathbf{x}_{L,2}]}, \dots), \quad (4)$$

where each $\mathbf{Y}_{L,i[\mathbf{x}_{L,i}]}$ corresponds to the potential outcomes of the variables $\mathbf{Y}_{L,i}$ under the intervention $\mathbf{X}_{L,i} = \mathbf{x}_{L,i}$. Each $\mathbf{Y}_{L,i}$ and $\mathbf{X}_{L,i}$ must be unions of clusters from \mathbb{C} (i.e. $\mathbf{Y}_{L,i} = \bigcup_{\mathbf{C} \in \mathbf{C}'} \mathbf{C}$ for some $\mathbf{C}' \subseteq \mathbb{C}$) such that $\tau(\mathbf{Y}_{L,i})$ and $\tau(\mathbf{X}_{L,i})$ are well-defined (i.e. $\tau(\mathbf{Y}_{L,i}) = (\bigwedge_{\mathbf{C} \in \mathbf{C}'} \tau_{\mathbf{C}}(\mathbf{C}))$). For instance, from Ex. 1, one term could be $\mathbf{Y}_{L,i} = \{D, C, F, P\} = \mathbf{C}_1 \cup \mathbf{C}_2$, but $\mathbf{Y}_{L,i} = \{D, C, F\}$ would be invalid since C and F are clustered together with P . For the high-level counterpart, denote

$$\mathbf{Y}_{H,*} = \tau(\mathbf{Y}_{L,*}) \quad (5)$$

$$= (\tau(\mathbf{Y}_{L,1[\tau(\mathbf{x}_{L,1})]}), \tau(\mathbf{Y}_{L,2[\tau(\mathbf{x}_{L,2})]}), \dots). \quad (6)$$

For any value $\mathbf{y}_{H,*} \in \mathcal{D}_{\mathbf{Y}_{H,*}}$, denote

$$\mathcal{D}_{\mathbf{Y}_{L,*}}(\mathbf{y}_{H,*}) = \{\mathbf{y}_{L,*} : \mathbf{y}_{L,*} \in \mathcal{D}_{\mathbf{Y}_{L,*}}, \tau(\mathbf{y}_{L,*}) = \mathbf{y}_{H,*}\}, \quad (7)$$

that is, the set of all values $\mathbf{y}_{L,*}$ such that $\tau(\mathbf{y}_{L,*}) = \mathbf{y}_{H,*}$. Considering again Ex. 1, if $\mathbf{y}_{H,*}$ refers to the value of $Z = 20$, then $\mathcal{D}_{\mathbf{Y}_{L,*}}(\mathbf{y}_{H,*})$ is the set of all tuples $(c, f, p) \in \mathcal{D}_{\mathbf{C}_2}$ such that $\tau(C = c, F = f, P = p) = (Z = 20)$ (i.e., $4c + 9f + 4p = 20$).

We can now define a notion of consistency relating low level counterfactual quantities to high level counterparts.

Definition 7 (Q - τ Consistency). Let \mathcal{M}_L and \mathcal{M}_H be SCMs defined over variables \mathbf{V}_L and \mathbf{V}_H , respectively. Let $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ be a constructive abstraction function w.r.t. clusters \mathbb{C} and \mathbb{D} . Let

$$Q = \sum_{\mathbf{y}_{L,*} \in \mathcal{D}_{\mathbf{Y}_{L,*}}(\mathbf{y}_{H,*})} P(\mathbf{Y}_{L,*} = \mathbf{y}_{L,*}) \quad (8)$$

be a low-level Layer 3 quantity of interest (for some $\mathbf{y}_{H,*} \in \mathcal{D}_{\mathbf{Y}_{H,*}}$), as expressed in Eq. 4, and let

$$\tau(Q) = P(\mathbf{Y}_{H,*} = \mathbf{y}_{H,*}) \quad (9)$$

be its high level counterpart, as expressed in Eq. 6. We say that \mathcal{M}_H is Q - τ consistent with \mathcal{M}_L if

$$\begin{aligned} & \sum_{\mathbf{y}_{L,*} \in \mathcal{D}_{\mathbf{Y}_{L,*}}(\mathbf{y}_{H,*})} P^{\mathcal{M}_L}(\mathbf{Y}_{L,*} = \mathbf{y}_{L,*}) \\ &= P^{\mathcal{M}_H}(\mathbf{Y}_{H,*} = \mathbf{y}_{H,*}), \end{aligned} \quad (10)$$

that is, the value of Q induced by \mathcal{M}_L is equal to the value of $\tau(Q)$ induced by \mathcal{M}_H ⁸. Furthermore, if \mathcal{M}_H is Q - τ consistent with \mathcal{M}_L for all $Q \in \mathcal{L}_i(\mathcal{M}_L)$ of the form of Eq. 8, then \mathcal{M}_H is said to be \mathcal{L}_i - τ consistent with \mathcal{M}_L . ■

In words, suppose Q is a quantity from \mathbf{V}_L in the form of Eq. 8. That is, it is a counterfactual quantity such that the variables of each term $\mathbf{Y}_{L,i}$ and each intervention $\mathbf{X}_{L,i}$ are unions of clusters in \mathbb{C} , and it is summed over values of $\mathbf{y}_{L,i}$ that map to one specific set of high level variables $\mathbf{y}_{H,i}$. Then, a query of this form has a counterpart $\tau(Q)$, obtained by applying τ on each term, shown in Eq. 9. We say that \mathcal{M}_H is Q - τ consistent with \mathcal{M}_L if the value of $\tau(Q)$, computed from \mathcal{M}_H , is equal to the value of Q computed from \mathcal{M}_L (i.e. Eq. 10 holds). Note that Def. 7 naturally applies to the \mathcal{L}_2 case (i.e. all $\mathbf{x}_{L,i}$ are identical) and the \mathcal{L}_1 case (i.e. all $\mathbf{X}_{L,i} = \emptyset$).

Def. 7 delineates the formal connection between quantities of \mathcal{M}_L and \mathcal{M}_H . Intuitively, \mathcal{M}_H can only be viewed as an abstraction of \mathcal{M}_L for the quantities in which they are τ -consistent. Consider the following example to ground the discussion.

Example 3. Consider a study on a new cancer drug. The drug is given in two doses, and patients take the second dose a month after the first dose. The variables observed are whether the individual takes the first dose (A), whether they take the second dose (B), whether they recover (Y), and whether they come from a wealthy background (R) and therefore have better nutrition and medical care. The SCM \mathcal{M}_L is as follows:

$$\mathbf{U}_L = \{U_{RY}, U_A, U_B, U_Y\} \quad (11)$$

$$\mathbf{V}_L = \{R, A, B, Y\} \quad (12)$$

$$\mathcal{F}_L = \begin{cases} R \leftarrow f_R^L(u_{RY}) = u_{RY} \\ A \leftarrow f_A^L(r, u_A) = r \oplus u_A \\ B \leftarrow f_B^L(r, a, u_B) = (r \wedge a) \oplus u_B \\ Y \leftarrow f_Y^L(a, b, u_{RY}, u_Y) \\ \quad = ((a \wedge b) \wedge u_{RY}) \oplus u_Y \end{cases} \quad (13)$$

$$P(\mathbf{U}_L) = \begin{cases} P(U_{RY} = 1) = 0.5 \\ P(U_A = 1) = P(U_B = 1) = 0.2 \\ P(U_Y = 1) = 0.1 \end{cases} \quad (14)$$

In words, people are more likely to take both doses if they are rich, and most people will only take the second dose if they have already taken the first dose. Also, people who take both doses are more likely to recover, but only if they came from a high socioeconomic background ($U_{RY} = 1$). The values computed from \mathcal{M}_L are shown in Figure 3.

The rows of the figure can be used to compute quantities of the PCH from \mathcal{M}_L . Denote p_i as the probability of the i th row of the table. Then, for example, the quantity $P(Y = 1 \mid A = 1, B = 1)$, or the probability that someone recovers given that they took both doses of the drug is

$$P^{\mathcal{M}_L}(Y = 1 \mid A = 1, B = 1)$$

⁸Note that the equality in Eq. 10 is consistent with the push-forward measure through τ .

	U_{RY}	U_A	U_B	U_Y	R	A	B	Y	$Y_{A=0,B=0}$	$Y_{A=1,B=1}$	P
0	0	0	0	0	0	0	0	0	0	0	$p_0 = 0.288$
1	0	0	0	1	0	0	0	1	1	1	$p_1 = 0.032$
2	0	0	1	0	0	0	1	0	0	0	$p_2 = 0.072$
3	0	0	1	1	0	0	1	1	1	1	$p_3 = 0.008$
4	0	1	0	0	0	1	0	0	0	0	$p_4 = 0.072$
5	0	1	0	1	0	1	0	1	1	1	$p_5 = 0.008$
6	0	1	1	0	0	1	1	0	0	0	$p_6 = 0.018$
7	0	1	1	1	0	1	1	1	1	1	$p_7 = 0.002$
8	1	0	0	0	1	1	1	1	0	1	$p_8 = 0.288$
9	1	0	0	1	1	1	1	0	1	0	$p_9 = 0.032$
10	1	0	1	0	1	1	0	0	0	1	$p_{10} = 0.072$
11	1	0	1	1	1	1	0	1	1	0	$p_{11} = 0.008$
12	1	1	0	0	1	0	0	0	0	1	$p_{12} = 0.072$
13	1	1	0	1	1	0	0	1	1	0	$p_{13} = 0.008$
14	1	1	1	0	1	0	1	0	0	1	$p_{14} = 0.018$
15	1	1	1	1	1	0	1	1	1	0	$p_{15} = 0.002$

Figure 3: Values computed from \mathcal{M}_L in Example 3.

$$\begin{aligned}
&= \frac{P^{\mathcal{M}_L}(Y = 1, A = 1, B = 1)}{P^{\mathcal{M}_L}(A = 1, B = 1)} \\
&= \frac{p_7 + p_8}{p_6 + p_7 + p_8 + p_9} \\
&= \frac{0.002 + 0.288}{0.018 + 0.002 + 0.288 + 0.032} \\
&\approx 0.853. \tag{15}
\end{aligned}$$

The causal quantity $P(Y_{A=1,B=1} = 1)$, or the probability that someone recovers when forced to take both doses of the drug, can be computed as

$$\begin{aligned}
&P^{\mathcal{M}_L}(Y_{A=1,B=1} = 1) \\
&= p_1 + p_3 + p_5 + p_7 + p_8 + p_{10} + p_{12} + p_{14} \\
&= 0.032 + 0.008 + 0.008 + 0.002 \\
&+ 0.288 + 0.072 + 0.072 + 0.018 \\
&= 0.5. \tag{16}
\end{aligned}$$

Indeed, one may be misled to think that the drug is extremely effective when only looking at the conditional quantity in Eq. 15, as opposed to the causal effect, as in Eq. 16. In reality, the causal effect of the drug is not as high.

Suppose the researchers decide that this much detail in the study is unnecessary, and they consider working in a more abstract model. One way to simplify the model is to reduce the amount of variables. Perhaps they decide that wealth (R) is irrelevant and can be abstracted away, and the two doses (A and B) can simply be abstracted into one variable, treatment (X). This can be represented using the intervariable clusters $\mathbb{C} = \{\mathbb{C}_1 = \{A, B\}, \mathbb{C}_2 = \{Y\}\}$, where A and B are in the same cluster, Y is in a separate cluster, and R is not included.

Further, the treatment X is only considered complete if both doses are taken, so we can further perform an intravari-able clustering, where the domains are:

$$\mathbb{D}_{\mathbb{C}_1} = \begin{cases} x_0 = \{(A = 0, B = 0), (A = 0, B = 1), \\ (A = 1, B = 0)\} \\ x_1 = \{(A = 1, B = 1)\} \end{cases} \tag{17}$$

We denote X as the higher level variable corresponding to \mathbb{C}_1 , and we define its domain $\mathcal{D}_X = \{0, 1\}$ to be binary, where x_0 corresponds to 0 and x_1 corresponds to 1. We leave Y as is in the lower level space.

We can then define the constructive abstraction function τ based on \mathbb{C} and \mathbb{D} , where $\mathbf{V}_H = \{X, Y\}$. For example,

$$\tau(R = 1, A = 0, B = 1, Y = 1) = (X = 0, Y = 1), \tag{18}$$

and

$$\tau(R = 1, A = 1, B = 1, Y = 0) = (X = 1, Y = 0). \tag{19}$$

Now define \mathcal{M}_H over the newly defined $\mathbf{V}_H = \tau(\mathbf{V}_L)$ as follows.

$$\mathbf{U}_H = \{U_X, U_{Y0}, U_{Y1}\} \tag{20}$$

$$\mathbf{V}_H = \{X, Y\} \tag{21}$$

$$\mathcal{F}_H = \begin{cases} X \leftarrow f_X^H(u_X) = u_X \\ Y \leftarrow f_Y^H(x, u_{Y0}, u_{Y1}) = \begin{cases} u_{Y0} & x = 0 \\ u_{Y1} & x = 1 \end{cases} \end{cases} \tag{22}$$

$$P(\mathbf{U}_H) = \begin{cases} P(U_X = 1) = 0.34 \\ P(U_{Y0} = 1) = 0.1 \\ P(U_{Y1} = 1) = 0.852941 \end{cases} \tag{23}$$

Interestingly, note that $P^{\mathcal{M}_H}(Y = 1 \mid X = 1) = P(U_{Y1} = 1) \approx 0.853$, which is equal to $P^{\mathcal{M}_L}(Y = 1 \mid A = 1, B = 1)$ computed in Eq. 15. In fact, if $Q = P(Y = 1, A = 1, B = 1)$, then the corresponding $\tau(Q)$ from Def. 7 is $P(Y = 1, X = 1)$, since $\tau(Y = 1, A = 1, B = 1) = (Y = 1, X = 1)$. Since they are equal, we would say that \mathcal{M}_H is Q - τ consistent with \mathcal{M}_L .

Now suppose $Q' = P(Y_{A=1,B=1} = 1)$. The corresponding $\tau(Q')$ would be $P(Y_{X=1} = 1)$. However, note that $P^{\mathcal{M}_H}(Y_{X=1} = 1) = P(U_{Y1} = 1) \approx 0.853$, which is not equal to $P^{\mathcal{M}_L}(Y_{A=1,B=1} = 1) = 0.5$ computed from Eq. 16. Then, \mathcal{M}_H is not Q' - τ consistent with \mathcal{M}_L .

It turns out that \mathcal{M}_H is Q - τ consistent with \mathcal{M}_L for every $Q \in \mathcal{L}_1$, making \mathcal{M}_H \mathcal{L}_1 - τ consistent with \mathcal{M}_L . On the other hand, this is not the case for \mathcal{L}_2 , the interventional layer. In fact, it seems that \mathcal{M}_H is equating correlation with causation and fails to capture the nuances of interventions in \mathcal{M}_L . Still, such a model could be useful if the queries of interest are on \mathcal{L}_1 . One could argue that \mathcal{M}_H is a suitable abstraction of \mathcal{M}_L on Layer 1, but not on Layer 2. The concept of Q - τ consistency allows us to define “partial” abstractions based on the specific quantities of the PCH that match. ■

It turns out that when \mathcal{M}_H is Q - τ consistent with \mathcal{M}_L on all three layers of the PCH (i.e. \mathcal{L}_3 - τ consistent), then \mathcal{M}_H can be considered an abstraction of \mathcal{M}_L on the SCM-level, which coincides with the definition of constructive τ -abstractions (Def. 17 from App. B) from Beckers and Halpern (2019), shown below.

Proposition 1 (Abstraction Connection). *Let $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ be a constructive abstraction function (Def. 6). \mathcal{M}_H is \mathcal{L}_3 - τ consistent (Def. 7) with \mathcal{M}_L if and only if there exists SCMs \mathcal{M}'_L and \mathcal{M}'_H s.t. $\mathcal{L}_3(\mathcal{M}'_L) = \mathcal{L}_3(\mathcal{M}_L)$, $\mathcal{L}_3(\mathcal{M}'_H) = \mathcal{L}_3(\mathcal{M}_H)$, and \mathcal{M}'_H is a constructive τ -abstraction of \mathcal{M}'_L .* ■

All proofs are provided in Appendix A. This proposition provides the connection between the abstractions defined in this work and established definitions from previous works⁹.

2.3 Algorithmic Abstraction Construction

With the abstraction function τ defined, the notion of Q - τ consistency allows for comparisons of distributions between the low level model \mathcal{M}_L and the abstraction \mathcal{M}_H . Still, it would be desirable to be able to systematically construct \mathcal{M}_H given \mathcal{M}_L and τ such that \mathcal{M}_H is Q - τ consistent with \mathcal{M}_L for as many queries Q as possible. Moving in this direction, we first note that as a subtlety, for some cases of \mathcal{M}_L , there are certain choices of \mathbb{C} and \mathbb{D} (and corresponding τ) for which Q - τ consistency (for some queries Q) is impossible to achieve in any choice of \mathcal{M}_H . This impossibility is illustrated in the following example, inspired by Spirtes and Scheines (2004).

Example 4. Consider a study that aims to understand the effects of diet on heart disease. Having a poor diet (X) is known to cause heart disease (Y) depending on its cholesterol content. Cholesterol comes in two forms, called high-density and low-density lipoproteins (HDL and LDL, respectively). The HDL is believed to lower heart disease risk while LDL increases it (Steinberg 2007; Truswell 2010). Suppose the study is simplified to binary variables, and the true model

⁹Note that one subtlety of this result is that it is not \mathcal{M}_H that is directly a constructive τ -abstraction of \mathcal{M}_L , but rather their \mathcal{L}_3 -equivalent counterparts, \mathcal{M}'_H and \mathcal{M}'_L . Indeed, the definition of constructive τ -abstractions is stronger than \mathcal{L}_3 - τ consistency (see proof for more details), but in tasks where we are only concerned with the layers of the PCH, this distinction is inconsequential.

\mathcal{M}_L is:

$$\mathbf{U}_L = \{U_X, U_{C1}, U_{C2}, U_Y\} \quad (24)$$

$$\mathbf{V}_L = \{X, HDL, LDL, Y\} \quad (25)$$

$$\mathcal{F}_L = \begin{cases} X \leftarrow f_X^L(u_X) = u_X \\ HDL \leftarrow f_{HDL}^L(x, u_{C1}) = x \oplus u_{C1} \\ LDL \leftarrow f_{LDL}^L(x, u_{C2}) = x \oplus u_{C2} \\ Y \leftarrow f_Y^L(hdl, ldl, u_Y) = (ldl \wedge \neg hdl) \oplus u_Y \end{cases} \quad (26)$$

$$P(\mathbf{U}_L) = \begin{cases} P(U_X = 1) = 0.5 \\ P(U_{C1} = 1) = 0.1 \\ P(U_{C2} = 1) = 0.1 \\ P(U_Y = 1) = 0.1 \end{cases} \quad (27)$$

As \mathcal{M}_L indicates, a person is more likely to get heart disease if their diet consists of high LDL levels but low HDL levels. For example, note that $P^{\mathcal{M}_L}(Y_{LDL=1, HDL=0} = 1) = 0.9$ while $P^{\mathcal{M}_L}(Y_{LDL=0, HDL=1} = 1) = 0.1$.

Now, suppose a data scientist decides to abstract HDL and LDL together into a variable called “total cholesterol” (TC). Say that TC is defined as

$$TC = HDL + LDL. \quad (28)$$

In fact, this leads to a choice of intervariable clusters

$$\mathbb{C} = \{\mathbf{C}_1 = \{X\}, \mathbf{C}_2 = \{HDL, LDL\}, \mathbf{C}_3 = \{Y\}\}, \quad (29)$$

and then for intravariabile clusters, they would choose

$$\mathbb{D}_{\mathbf{C}_2} = \begin{cases} tc_0 = \{(HDL = 0, LDL = 0)\} \\ tc_1 = \{(HDL = 0, LDL = 1), \\ (HDL = 1, LDL = 0)\} \\ tc_2 = \{(HDL = 1, LDL = 1)\}. \end{cases} \quad (30)$$

For the other clusters, simply use the same variables. Let τ be the constructive abstraction function defined with this choice of \mathbb{C} and \mathbb{D} (i.e. $\tau_{\mathbf{C}_2}(hdl, ldl) = hdl + ldl$).

An issue arises due to the grouping of values ($HDL = 0, LDL = 1$) and ($HDL = 1, LDL = 0$) into the same intravariabile cluster. To witness, note that $\tau_{\mathbf{C}_1}(HDL = 0, LDL = 1) = \tau_{\mathbf{C}_2}(HDL = 1, LDL = 0) = (TC = 1)$. Now, consider two queries $Q_1 = P(Y_{HDL=0, LDL=1} = 1)$ and $Q_2 = P(Y_{HDL=1, LDL=0} = 1)$, and observe that

$$P^{\mathcal{M}_L}(Y_{HDL=0, LDL=1} = 1) = P(U_Y = 0) = 0.9, \quad (31)$$

$$P^{\mathcal{M}_L}(Y_{HDL=1, LDL=0} = 1) = P(U_Y = 1) = 0.1. \quad (32)$$

However, since $\tau_{\mathbf{C}_1}(HDL = 0, LDL = 1) = \tau_{\mathbf{C}_2}(HDL = 1, LDL = 0) = (TC = 1)$, both Q_1 and Q_2 have the same high-level counterpart. That is, $\tau(Q_1) = \tau(Q_2) = P(Y_{TC=1} = 1)$. No choice of \mathcal{M}_H over \mathbf{V}_H can be both Q_1 - τ consistent and Q_2 - τ consistent with \mathcal{M}_L because $P^{\mathcal{M}_H}(Y_{TC=1} = 1)$ cannot both be equal to 0.9 and 0.1. ■

Intuitively, Ex. 4 shows two values that cannot be grouped into the same intravariabile cluster because the function f_Y^L (from Eq. 26) produces different results depending on which

value is used. Grouping the two values in the same cluster would imply that the two values are “equivalent” and hence, there would be an inevitable loss of information. Indeed, real-world studies that consider total cholesterol instead of separating it into LDL and HDL often have conflicting results, indicating an invalid abstraction. This phenomenon can be described formally through the following condition.

Definition 8 (Abstract Invariance Condition (AIC)). Let $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle$ be an SCM and $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ be a constructive abstraction function relative to \mathbb{C} and \mathbb{D} . The SCM \mathcal{M}_L is said to satisfy the abstract invariance condition (AIC, for short) with respect to τ if, for all $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{D}_{\mathbf{V}_L}$ such that $\tau(\mathbf{v}_1) = \tau(\mathbf{v}_2)$, $\forall \mathbf{u} \in \mathcal{D}_{\mathbf{U}_L}, \mathbf{C}_i \in \mathbb{C}$, the following holds:

$$\begin{aligned} & \tau_{\mathbf{C}_i} \left(\left(f_V^L(\mathbf{pa}_V^{(1)}, \mathbf{u}_V) : V \in \mathbf{C}_i \right) \right) \\ &= \tau_{\mathbf{C}_i} \left(\left(f_V^L(\mathbf{pa}_V^{(2)}, \mathbf{u}_V) : V \in \mathbf{C}_i \right) \right), \end{aligned} \quad (33)$$

where $\mathbf{pa}_V^{(1)}$ and $\mathbf{pa}_V^{(2)}$ are the values corresponding to \mathbf{v}_1 and \mathbf{v}_2 . Then, $\widetilde{\mathbf{pa}}_V$ is used to denote any arbitrary value s.t. $\tau(\widetilde{\mathbf{pa}}_V) = \tau(\mathbf{pa}_V^{(1)}) = \tau(\mathbf{pa}_V^{(2)})$. ■

In words, the AIC enforces that if two low level values $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{D}_{\mathbf{V}_L}$ map to the same high level value (i.e. $\tau(\mathbf{v}_1) = \tau(\mathbf{v}_2)$), then for each cluster $\mathbf{C}_i \in \mathbb{C}$, the functions of those clusters should map to the same value regardless of \mathbf{U}_L (i.e. the outputs of $f_V^L(\mathbf{pa}_V^{(1)}, \mathbf{u}_V)$ for each $V \in \mathbf{C}_i$ should map to the same result as the outputs of $f_V^L(\mathbf{pa}_V^{(2)}, \mathbf{u}_V)$ when passed through $\tau_{\mathbf{C}_i}$). Intuitively, this implies that two values in the same intravariabile cluster have the same functional effect in the higher level setting.

In Ex. 4, the AIC is not satisfied since $(HDL = 0, LDL = 1)$ and $(HDL = 1, LDL = 0)$ cannot be grouped into the same intravariabile cluster. As established, $\tau_{\mathbf{C}_1}(HDL = 0, LDL = 1) = \tau_{\mathbf{C}_2}(HDL = 1, LDL = 0) = 1$. However, observing f_Y^L (from Eq. 26) given these inputs, note that $f_Y^L(HDL = 0, LDL = 1, u_Y) = \neg u_Y$, while $f_Y^L(HDL = 1, LDL = 0, u_Y) = u_Y$, which give opposite results for either choice of u_Y .

In contrast, we consider a different abstraction that does satisfy the AIC next.

Example 5 (Example 4 continued). Consider a different choice of clusters that does not violate the AIC. Although somewhat unintuitive, we can actually abstract the cholesterol values in the given \mathcal{M}_L by taking their difference instead of their sum!¹⁰ For example, define $Z = LDL - HDL$. In terms of clusters we would keep the same intervariable clusters $\mathbb{C} = \{\mathbf{C}_1 = \{X\}, \mathbf{C}_2 = \{HDL, LDL\}, \mathbf{C}_3 = \{Y\}\}$, but for intravariabile clusters, we choose $\mathbb{D}_{\mathbf{C}_2} = \{z_{-1} = \{(HDL = 1, LDL = 0)\}, z_0 = \{(HDL = 0, LDL = 0), (HDL = 1, LDL = 1)\}, z_1 = \{(HDL = 0, LDL = 1)\}\}$. Define τ as the constructive abstraction function defined over \mathbb{C} and \mathbb{D} (i.e. $\tau_{\mathbf{C}_2}(hdl, ldl) = ldl - hdl$).

¹⁰This is true given the simple definition of \mathcal{M}_L in this example. In more complex descriptions, such as in cases where the variables are continuous, the clusters would have to be chosen more carefully to avoid violations of the AIC.

Algorithm 1: Constructing \mathcal{M}_H from \mathcal{M}_L .

Input : SCM $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle$,
admissible inter/intravariabile clusters \mathbb{C} and
 \mathbb{D} satisfying abstract invariance condition
Output : SCM \mathcal{M}_H and $\tau : \mathcal{D}_{\mathbf{V}_H} \rightarrow \mathcal{D}_{\mathbf{V}_L}$ s.t. \mathcal{M}_H
is \mathcal{L}_3 - τ consistent with \mathcal{M}_L

```

1  $\mathbf{U}_H \leftarrow \mathbf{U}_L, P(\mathbf{U}_H) \leftarrow P(\mathbf{U}_L)$ 
2  $\mathbf{V}_H \leftarrow \mathbb{C}, \mathcal{D}_{\mathbf{V}_H} \leftarrow \mathbb{D}$ 
3  $\tau \leftarrow \text{AbsFunc}(\mathbb{C}, \mathbb{D})$  // from Def. 6
4 for  $\mathbf{C}_i \in \mathbb{C}$  do
5    $\lfloor f_i^H \leftarrow \tau(f_V^L(\widetilde{\mathbf{pa}}_V, \mathbf{u}_V) : V \in \mathbf{C}_i)$ 
6  $\mathcal{F}_H \leftarrow \{f_i^H : \mathbf{C}_i \in \mathbb{C}\}$ 
7 return  $\tau, \mathcal{M}_H = \langle \mathbf{U}_H, \mathbf{V}_H, \mathcal{F}_H, P(\mathbf{U}_H) \rangle$ 

```

The difference now is that instead of clustering together the values $(HDL = 1, LDL = 0)$ and $(HDL = 0, LDL = 1)$, it is the values $(HDL = 0, LDL = 0)$ and $(HDL = 1, LDL = 1)$ that are clustered together. When looking at f_Y^L (Eq. 26) under these two values, we see that $f_Y^L(HDL = 0, LDL = 0, u_Y) = f_Y^L(HDL = 1, LDL = 1, u_Y) = u_Y$ for any choice of u_Y , satisfying the AIC.

Intuitively, f_Y no longer changes behavior between these two values, or in other words, f_Y is *invariant* between these two values. With no other downstream variables to consider, this implies that these two values are functionally identical in the model and can be abstracted together into a single value without loss of information. ■

It turns out that the AIC describes precisely when an appropriate \mathcal{M}_H exists as an abstraction of the low level model \mathcal{M}_L , as shown by the following result.

Proposition 2 (Abstraction Conditions). *For any SCM \mathcal{M}_L and constructive abstraction function τ relative to \mathbb{C} and \mathbb{D} , there exists an SCM \mathcal{M}_H over variables $\mathbf{V}_H = \tau(\mathbf{V}_L)$ such that \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L if and only if there exists \mathcal{M}'_L such that $\mathcal{L}_3(\mathcal{M}_L) = \mathcal{L}_3(\mathcal{M}'_L)$ and \mathcal{M}'_L satisfies the abstract invariance condition with respect to τ .* ■

This critical property guarantees the existence of a higher level SCM \mathcal{M}_H such that \mathcal{L}_3 - τ consistency holds, so we will assume that the AIC holds for the rest of this work. Still, see App. D.2 for further discussion on its implications and for possible relaxations in cases where \mathcal{L}_3 - τ consistency is not required.

With the notion of abstractions well-defined, we study how \mathcal{M}_H can be obtained from \mathcal{M}_L . Interestingly, when given the admissible clusterings \mathbb{C} and \mathbb{D} , the procedure for recovering τ and converting \mathcal{M}_L to \mathcal{M}_H can be done as shown in Alg. 1. Intuitively, one can obtain an abstraction \mathcal{M}_H of \mathcal{M}_L by first constructing the abstraction function τ using the clusterings \mathbb{C} and \mathbb{D} (lines 2-3), followed by designing the functions of \mathcal{M}_H to wrap the original functions of \mathcal{M}_L with τ (lines 4-6). This can be verified using the following result.

Proposition 3. *Let τ and \mathcal{M}_H be the function and SCM obtained from running Alg. 1 on inputs $\mathcal{M}_L, \mathbb{C}$, and \mathbb{D} . Then, \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L .* ■

See below for an example of running Alg. 1

Example 6 (Example 3 continued). We will run Alg. 1 on \mathcal{M}_L and clusters \mathbb{C} and \mathbb{D} described earlier in Ex. 3. Following the algorithm, we first set $\mathbf{U}_H = \mathbf{U}_L$ and $P(\mathbf{U}_H) = P(\mathbf{U}_L)$. We construct τ via Def. 6, as shown in the earlier example. Then, we can compute the function f_X^H as follows.

$$\begin{aligned} f_X^H(u_{RY}, u_A, u_B) &= \tau_{\mathbb{C}_1}(f_A^L(r, u_A), f_B^L(r, u_A, u_B)) \\ &= \tau_{\mathbb{C}_1}(r \oplus u_A, (r \wedge (r \oplus u_A)) \oplus u_B) \\ &= \tau_{\mathbb{C}_1}(u_{RY} \oplus u_A, (u_{RY} \wedge (u_{RY} \oplus u_A)) \oplus u_B) \end{aligned} \quad (34)$$

For f_Y^H , denote \tilde{a} and \tilde{b} as an arbitrary setting of A and B such that $\tau(\tilde{a}, \tilde{b}) = x$, as indicated in line 5 of the algorithm.

$$\begin{aligned} f_Y^H(x, u_{RY}, u_Y) &= \tau_{\mathbb{C}_2}(f_Y^L(\tilde{a}, \tilde{b}, u_{RY}, u_Y)) \\ &= \tau_{\mathbb{C}_2}(((\tilde{a} \wedge \tilde{b}) \wedge u_{RY}) \oplus u_Y) \\ &= \tau_{\mathbb{C}_2}((x \wedge u_{RY}) \oplus u_Y) \end{aligned} \quad (35)$$

Putting everything together, we obtain \mathcal{M}_H as follows.

$$\mathbf{U}_H = \mathbf{U}_L = \{U_{RY}, U_A, U_B, U_Y\} \quad (36)$$

$$\mathbf{V}_H = \{X, Y\} \quad (37)$$

$$\mathcal{F}_H = \{ \quad (38)$$

$$\begin{aligned} X &\leftarrow f_X^H(u_{RY}, u_A, u_B) \\ &= \tau_{\mathbb{C}_1}(u_{RY} \oplus u_A, (u_{RY} \wedge (u_{RY} \oplus u_A)) \oplus u_B) \end{aligned}$$

$$Y \leftarrow f_Y^H(x, u_{RY}, u_Y) = \tau_{\mathbb{C}_2}((x \wedge u_{RY}) \oplus u_Y)$$

$$P(\mathbf{U}_H) = P(\mathbf{U}_L) \quad (39)$$

It is not difficult to see that \mathcal{M}_H is $\mathcal{L}_{3-\tau}$ consistent with \mathcal{M}_L . As an example, note that $P^{\mathcal{M}_H}(Y_{X=1} = 1) = P(U_{RY} = 1, U_Y = 0) + P(U_{RY} = 0, U_Y = 1) = (0.5)(0.9) + (0.5)(0.1) = 0.5$, which matches $P^{\mathcal{M}_L}(Y_{A=1, B=1} = 1)$ from Eq. 16. ■

Ex. 6 shows how Alg. 1 can be used to systematically obtain an abstraction \mathcal{M}_H of the low-level model \mathcal{M}_L , so long as \mathcal{M}_L is provided alongside the clusters \mathbb{C} and \mathbb{D} . Since \mathcal{M}_L is almost never available in practice, the following sections show how this requirement can be relaxed.

3 Inferences Across Abstractions

As demonstrated by Alg. 1, converting a low level model \mathcal{M}_L to a high level model \mathcal{M}_H is somewhat immediate when given full observability of the underlying SCM \mathcal{M}_L . However, in real applications, it is rarely the case that the full specification of \mathcal{M}_L is known. Typically, one will only be given partial information of \mathcal{M}_L in the form of data, such as samples of the observational distribution $P(\mathbf{V}_L)$. The question we investigate in this section is: is it still possible to “learn” some \mathcal{M}_H given the observed data?

We first note the impossibility result described by the Causal Hierarchy Theorem (CHT) (Bareinboim et al. 2022, Thm. 1), which states that a model trained to match another SCM on lower layers of the causal hierarchy (e.g. \mathcal{L}_1) will

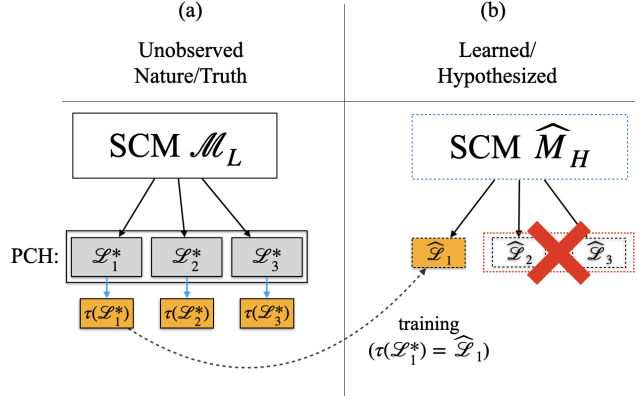


Figure 4: Illustration of the Abstract CHT. Without additional information, a high-level model $\widehat{\mathcal{M}}_H$ trained to be $\mathcal{L}_{1-\tau}$ consistent with \mathcal{M}_L is not guaranteed to be \mathcal{L}_2 or $\mathcal{L}_{3-\tau}$ consistent.

likely not match on higher layers (e.g. \mathcal{L}_2 or \mathcal{L}_3). Naturally, the same is true when it comes to inferring causal quantities across abstractions. One may be tempted to believe that \mathcal{M}_H can be learned given \mathcal{L}_1 data from \mathcal{M}_L by instantiating some expressive parametric model $\widehat{\mathcal{M}}_H$ on \mathbf{V}_H , and then training $\widehat{\mathcal{M}}_H$ on $P(\mathbf{V}_H) = P(\tau(\mathbf{V}_L))$ such that $\widehat{\mathcal{M}}_H$ is $\mathcal{L}_{1-\tau}$ consistent with \mathcal{M}_L . Unfortunately, this strategy will fail in general since even under perfect training, $\widehat{\mathcal{M}}_H$ is not guaranteed to be $\mathcal{L}_{2-\tau}$ (or $\mathcal{L}_{3-\tau}$) consistent with \mathcal{M}_L . This means that any causal quantities induced by $\widehat{\mathcal{M}}_H$ will likely bear no relationship with causal quantities induced by \mathcal{M}_L . This phenomenon is described by the following proposition.

Proposition 4 (Abstract Causal Hierarchy Theorem (Informal)). *Given constructive abstraction function $\tau : \mathcal{D}_{\mathbf{V}_H} \rightarrow \mathcal{D}_{\mathbf{V}_L}$, even if \mathcal{M}_H is \mathcal{L}_i - τ consistent with \mathcal{M}_L , \mathcal{M}_H will almost never be \mathcal{L}_j - τ consistent with \mathcal{M}_L for $j > i$. ■*

In words, matching across abstractions on lower layers does not guarantee the same will hold for higher layers. This idea is illustrated in Fig. 4. The left of the figure shows the unobserved true SCM \mathcal{M}_L , which induces distributions from the three layers of the PCH. Observational data from \mathcal{L}_1 is provided, and one may train a high-level model $\widehat{\mathcal{M}}_H$ (right) such that it is $\mathcal{L}_{1-\tau}$ consistent with \mathcal{M}_L . However, even if training is perfect, it is not guaranteed that $\widehat{\mathcal{M}}_H$ is \mathcal{L}_2 or $\mathcal{L}_{3-\tau}$ consistent with \mathcal{M}_L . That is, queries from the \mathcal{L}_2 or \mathcal{L}_3 distributions of $\widehat{\mathcal{M}}_H$ are not expected to match the equivalent queries in \mathcal{M}_L in general. (See further details in App. A.)

The consequence of this result is that causal assumptions will be necessary to make progress. In particular, the class of plausible models must be constrained through assumptions about the generating model. Given this necessity, one type of assumption prevalent throughout causal inference literature is the availability of a causal diagram (Pearl 1995), a graphical structure that qualitatively describes the functional relationships between variables in a non-parametric manner. This assumption is weaker than having the entire generat-

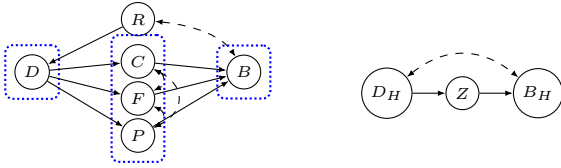


Figure 5: The causal diagram \mathcal{G} over variables \mathbf{V}_L for the nutrition study in Ex. 1 is on the left. Clusters $\mathbb{C} = \{D_H = \{D\}, Z = \{C, F, P\}, B_H = \{B\}\}$ are outlined in blue. The corresponding C-DAG $\mathcal{G}_\mathbb{C}$ is on the right.

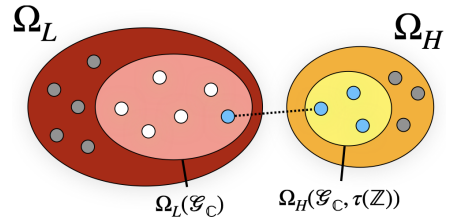
ing SCM, since it only encodes qualitative information of the functional dependences between exogenous and endogenous variables (as in Def. 2) rather than full detail of the generating mechanisms and exogenous distributions. Still, it has been shown that having the causal diagram allows certain inferences across layers, determined through the causal identification problem (Pearl 2000; Bareinboim and Pearl 2016).

In the context of abstractions however, specifying the causal diagram for the true model \mathcal{M}_L requires describing the relationships between every low-level variable in \mathbf{V}_L . This may be unrealistic in many practical settings since there are typically too many low-level variables (e.g. 128×128 pixels in an image) to expect a description of the relationship between every pair, and many of these relationships may not be well-defined in a causal manner. Instead, it may be more reasonable to specify a causal diagram over \mathbf{V}_H (or intervariable clusters \mathbb{C}). The amount of information required is reduced when $|\mathbf{V}_H| \ll |\mathbf{V}_L|$, and the causal relationships between variables may be more clear given that the higher-level variables tend to be more explainable. The causal diagram over \mathbf{V}_H can be viewed as a graphical abstraction of the causal diagram over \mathbf{V}_L . The relationship can be formalized through the concept of cluster causal diagrams (C-DAGs) (Anand et al. 2023), as described next.

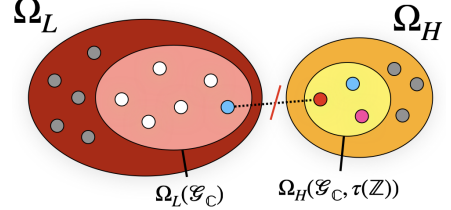
Definition 9 (Cluster Causal Diagram (C-DAG) (Anand et al. 2023, Def. 1)). Given a causal diagram $\mathcal{G} = \langle \mathbf{V}, \mathbf{E} \rangle$ and an admissible clustering $\mathbb{C} = \{\mathbb{C}_1, \dots, \mathbb{C}_k\}$ of \mathbf{V} , construct a graph $\mathcal{G}_\mathbb{C} = \langle \mathbb{C}, \mathbf{E}_\mathbb{C} \rangle$ over \mathbb{C} with a set of edges $\mathbf{E}_\mathbb{C}$ defined as follows:

1. A directed edge $\mathbb{C}_i \rightarrow \mathbb{C}_j$ is in $\mathbf{E}_\mathbb{C}$ if there exists some $V_i \in \mathbb{C}_i$ and $V_j \in \mathbb{C}_j$ such that $V_i \rightarrow V_j$ is an edge in \mathbf{E} .
2. A dashed bidirected edge $\mathbb{C}_i \leftrightarrow \mathbb{C}_j$ is in $\mathbf{E}_\mathbb{C}$ if there exists some $V_i \in \mathbb{C}_i$ and $V_j \in \mathbb{C}_j$ such that $V_i \leftrightarrow V_j$ is an edge in \mathbf{E} . ■

In words, the nodes of the C-DAG $\mathcal{G}_\mathbb{C}$ simply correspond to the clusters of \mathbb{C} , and edges connect clusters \mathbb{C}_i and \mathbb{C}_j if they connect some $V_i \in \mathbb{C}_i$ and $V_j \in \mathbb{C}_j$ in the original causal diagram \mathcal{G} . Interestingly, the C-DAG definition aligns with the concept of intervariable clusters, providing a way for encoding constraints in the smaller space of \mathbf{V}_H . Revisiting the nutrition study in Ex. 1, Fig. 5 shows the corresponding causal diagram \mathcal{G} (left) and the simpler C-DAG $\mathcal{G}_\mathbb{C}$ (right). With the constraints of $\mathcal{G}_\mathbb{C}$, we now introduce a notion of identification across abstractions to determine precisely which queries can be inferred.



(a) When Q is τ -ID from $\mathcal{G}_\mathbb{C}$ and \mathbb{Z} , all SCMs over \mathbf{V}_H that induce $\mathcal{G}_\mathbb{C}$ and is \mathbb{Z} - τ consistent with \mathcal{M}_L are also Q - τ consistent with \mathcal{M}_L for any choice of \mathcal{M}_L .



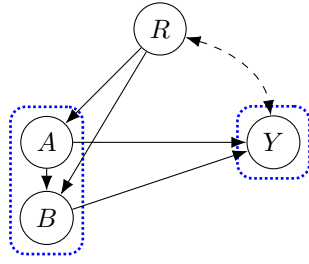

(b) When Q is not τ -ID from $\mathcal{G}_\mathbb{C}$ and \mathbb{Z} , then there exist choices of \mathcal{M}_L such that some SCM \mathcal{M}_H over \mathbf{V}_H that induces $\mathcal{G}_\mathbb{C}$ and is \mathbb{Z} - τ consistent with \mathcal{M}_L is not Q - τ consistent with \mathcal{M}_L .

Figure 6: Examples of τ -ID and τ -nonID cases. The space of models over \mathbf{V}_L , Ω_L , is shown in dark red (left), and the subspace that induces C-DAG $\mathcal{G}_\mathbb{C}$ is shown in light red. The blue dot within this space is an arbitrary choice of \mathcal{M}_L . The space of models over \mathbf{V}_H , Ω_H , is shown in dark yellow (right), and the subspace that induces the causal graph $\mathcal{G}_\mathbb{C}$ and is also \mathbb{Z} - τ consistent with \mathcal{M}_L is shown in light yellow.

Definition 10 (Abstract Identification). Let $\tau : \mathcal{D}_{\mathbf{V}_H} \rightarrow \mathcal{D}_{\mathbf{V}_L}$ be a constructive abstraction function. Consider C-DAG $\mathcal{G}_\mathbb{C}$, and let $\mathbb{Z} = \{P(\mathbf{V}_{L[\mathbf{z}_k]})\}_{k=1}^{\ell}$ be a collection of available interventional (or observational if $\mathbf{Z}_k = \emptyset$) distributions over \mathbf{V}_L . Let Ω_L and Ω_H be the space of SCMs defined over \mathbf{V}_L and \mathbf{V}_H , respectively, and let $\Omega_L(\mathcal{G}_\mathbb{C})$ and $\Omega_H(\mathcal{G}_\mathbb{C})$ be their corresponding subsets that induce C-DAG $\mathcal{G}_\mathbb{C}$. A query Q is said to be τ -ID from $\mathcal{G}_\mathbb{C}$ and \mathbb{Z} iff for every $\mathcal{M}_L \in \Omega_L(\mathcal{G}_\mathbb{C}), \mathcal{M}_H \in \Omega_H(\mathcal{G}_\mathbb{C})$ such that \mathcal{M}_H is \mathbb{Z} - τ consistent with \mathcal{M}_L , \mathcal{M}_H is also Q - τ consistent with \mathcal{M}_L . ■

This definition establishes a notion of identification between two different spaces of SCMs, Ω_L and Ω_H , that are connected through τ . The concept of τ -ID is illustrated in Fig. 6. In words, τ -identifiability implies that in every pair of SCMs \mathcal{M}_L over \mathbf{V}_L and \mathcal{M}_H over \mathbf{V}_H , “matching” in graph $\mathcal{G}_\mathbb{C}$ and data \mathbb{Z} implies a match in query Q . Since \mathcal{M}_L and \mathcal{M}_H are defined over different spaces of variables, the term “match” has some nuance. Specifically, “matching” in $\mathcal{G}_\mathbb{C}$ implies that $\mathcal{G}_\mathbb{C}$ is a C-DAG for \mathcal{M}_L and is a causal diagram for \mathcal{M}_H . “Matching” in \mathbb{Z} (resp. Q) implies that \mathcal{M}_H is \mathbb{Z} - τ consistent (resp. Q - τ consistent) with \mathcal{M}_L . As shown in Fig. 6a, τ -ID implies that for any choice of \mathcal{M}_L over \mathbf{V}_L that induces C-DAG $\mathcal{G}_\mathbb{C}$ (blue dot in light red space), all SCMs \mathcal{M}_H over \mathbf{V}_H that induce $\mathcal{G}_\mathbb{C}$ and are \mathbb{Z} - τ consistent with \mathcal{M}_L (light yellow space) are also Q - τ consistent with \mathcal{M}_L . As a consequence, $\tau(Q)$ can be evaluated from any of

R	A	B	Y	P
0	0	0	0	0.288
0	0	0	1	0.032
0	0	1	0	0.072
0	0	1	1	0.008
0	1	0	0	0.072
0	1	0	1	0.008
0	1	1	0	0.018
0	1	1	1	0.002
1	0	0	0	0.072
1	0	0	1	0.008
1	0	1	0	0.018
1	0	1	1	0.002
1	1	0	0	0.072
1	1	0	1	0.008
1	1	1	0	0.032
1	1	1	1	0.288

(a) $P^{\mathcal{M}_L}(\mathbf{V}_L)$.(b) Causal diagram \mathcal{G} . Intervariable clusters \mathbb{C} are grouped by the dashed-blue lines.(c) C-DAG \mathcal{G}_C .

\mathbf{V}_H term	Corresponding \mathbf{V}_L term	P
$P(X = 0, Y = 0)$	$\sum_r \sum_{a,b:(a,b) \neq (1,1)} P(R = r, A = a, B = b, Y = 0)$	0.594
$P(X = 0, Y = 1)$	$\sum_r \sum_{a,b:(a,b) \neq (1,1)} P(R = r, A = a, B = b, Y = 1)$	0.066
$P(X = 1, Y = 0)$	$\sum_r P(R = r, A = 1, B = 1, Y = 0)$	0.05
$P(X = 1, Y = 1)$	$\sum_r P(R = r, A = 1, B = 1, Y = 1)$	0.29

(d) $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$.Figure 7: Dataset quantities computed from \mathcal{M}_L and graphs for Example 7.

these choices of \mathcal{M}_H to compute Q .

On the other hand, τ -nonidentifiability implies that there exist a pair of models \mathcal{M}_L over \mathbf{V}_L and \mathcal{M}_H over \mathbf{V}_H such that \mathcal{M}_L and \mathcal{M}_H match in both \mathcal{G}_C and \mathbb{Z} yet still do not match in Q . As shown in Fig. 6b, τ -non-ID implies that there exists some \mathcal{M}_L over \mathbf{V}_L that induces C-DAG \mathcal{G}_C (blue dot in light red space) such that there is some \mathcal{M}_H over \mathbf{V}_H that induces \mathcal{G}_C and is \mathbb{Z} - τ consistent with \mathcal{M}_L (light yellow space) but is not Q - τ consistent with \mathcal{M}_L . This means that despite the constraints added through the C-DAG \mathcal{G}_C , there are still queries that cannot be inferred across τ due to nonidentifiability. This is more acute when there is a large amount of unobserved confounding. Consider the following example where this is the case.

Example 7 (Example 3 continued). The data scientist team may be interested in computing the causal effect of taking both stages of the drug on the recovery rate of the disease, $Q = P(Y_{A=1, B=1} = 1)$. However, \mathcal{M}_L is not observed, and instead, the observational data from $P(\mathbf{V}_L)$ ($\mathbb{Z} = \{P(\mathbf{V}_L)\}$) and the C-DAG \mathcal{G}_C from Fig. 7c are given. Is Q τ -ID from \mathbb{Z} and \mathcal{G}_C ? It turns out the answer is no. To witness, consider the following pair of models.

$$\mathcal{M}_1 = \begin{cases} \mathbf{U}_1 = \{U_{XY}, U_{Y0}, U_{Y1}\} \\ \mathbf{V}_H = \{X, Y\} \\ \mathcal{F}_1 = \begin{cases} X \leftarrow f_X^1(u_{XY}) = u_{XY} \\ Y \leftarrow f_Y^1(x, u_{XY}, u_{Y0}, u_{Y1}) \\ = \begin{cases} x \vee u_{Y0} & u_{XY} = 0 \\ u_{Y1} & u_{XY} = 1 \end{cases} \end{cases} \\ P(\mathbf{U}_1) = \begin{cases} P(U_{XY} = 1) = 0.34 \\ P(U_{Y0} = 1) = 0.1 \\ P(U_{Y1} = 1) = 0.852941 \end{cases} \end{cases} \quad (40)$$

$$\mathcal{M}_2 = \begin{cases} \mathbf{U}_2 = \{U_{XY}, U_{Y0}, U_{Y1}\} \\ \mathbf{V}_H = \{X, Y\} \\ \mathcal{F}_2 = \begin{cases} X \leftarrow f_X^2(u_{XY}) = u_{XY} \\ Y \leftarrow f_Y^2(x, u_{XY}, u_{Y0}, u_{Y1}) \\ = \begin{cases} \neg x \wedge u_{Y0} & u_{XY} = 0 \\ u_{Y1} & u_{XY} = 1 \end{cases} \end{cases} \\ P(\mathbf{U}_2) = \begin{cases} P(U_{XY} = 1) = 0.34 \\ P(U_{Y0} = 1) = 0.1 \\ P(U_{Y1} = 1) = 0.852941 \end{cases} \end{cases} \quad (41)$$

Every value of $P^{\mathcal{M}_L}(\mathbf{V}_L)$ can be computed using the table in Fig. 3, producing the table in Fig. 7a. The corresponding values of $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$ can then be computed as shown in Fig. 7d via Eq. 10. One can verify that both $P^{\mathcal{M}_1}(\mathbf{V}_H)$ and $P^{\mathcal{M}_2}(\mathbf{V}_H)$ match the values in Fig. 7d, implying that both \mathcal{M}_1 and \mathcal{M}_2 are $P(\mathbf{V}_L)$ - τ consistent with \mathcal{M}_L . Moreover, \mathcal{M}_1 and \mathcal{M}_2 also induce \mathcal{G}_C from Fig. 7c: f_Y takes X as input, and f_X and f_Y share U_{XY} as a confounding variable. Also, f_X does not contain Y as an input.

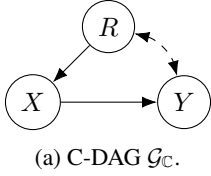
However, computing $\tau(Q) = P(Y_{X=1} = 1)$ leads to two different answers, i.e.:

$$\begin{aligned} P^{\mathcal{M}_1}(Y_{X=1} = 1) &= P(U_{XY} = 0) + P(U_{XY} = 1, U_{Y1} = 1) \\ &= 0.66 + (0.34)(0.852941) = 0.95, \end{aligned} \quad (42)$$

and

$$\begin{aligned} P^{\mathcal{M}_2}(Y_{X=1} = 1) &= P(U_{XY} = 1, U_{Y1} = 1) \\ &= (0.34)(0.852941) = 0.29. \end{aligned} \quad (43)$$

These values are not only different, but also neither are equal to the true value $P^{\mathcal{M}_L}(Y_{A=1, B=1} = 1) = 0.5$, as computed



\mathbf{V}_H term	Corresponding \mathbf{V}_L term	P
$P(R = 0, X = 0, Y = 0)$	$\sum_{a,b:(a,b) \neq (1,1)} P(R = 0, A = a, B = b, Y = 0)$	0.432
$P(R = 0, X = 0, Y = 1)$	$\sum_{a,b:(a,b) \neq (1,1)} P(R = 0, A = a, B = b, Y = 1)$	0.048
$P(R = 0, X = 1, Y = 0)$	$P(R = 0, A = 1, B = 1, Y = 0)$	0.018
$P(R = 0, X = 1, Y = 1)$	$P(R = 0, A = 1, B = 1, Y = 1)$	0.002
$P(R = 1, X = 0, Y = 0)$	$\sum_{a,b:(a,b) \neq (1,1)} P(R = 1, A = a, B = b, Y = 0)$	0.162
$P(R = 1, X = 0, Y = 1)$	$\sum_{a,b:(a,b) \neq (1,1)} P(R = 1, A = a, B = b, Y = 1)$	0.018
$P(R = 1, X = 1, Y = 0)$	$P(R = 1, A = 1, B = 1, Y = 0)$	0.032
$P(R = 1, X = 1, Y = 1)$	$P(R = 1, A = 1, B = 1, Y = 1)$	0.288

(b) $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$.

Figure 8: Updated items given the new choice of \mathbb{C} in Example 8

in Eq. 16. A scientist using model \mathcal{M}_1 may conclude that the treatment is extremely effective, while a scientist using model \mathcal{M}_2 may conclude the opposite: the treatment is not only ineffective, it is even harmful. The query is not τ -ID in this case, and no further inferences of this query should be made at this stage. ■

Now consider the following τ -ID example.

Example 8 (Example 3 continued). The data scientist team is studying causal inference and notes that Example 7 showed that the query $Q = P(Y_{A=1, B=1} = 1)$ is non-ID from the available data. However, suppose instead, a different set of clusters is used that includes R , i.e. $\mathbb{C} = \{\mathbb{C}_1 = \{A, B\}, \mathbb{C}_2 = \{Y\}, \mathbb{C}_3 = \{R\}\}$. The domain of R would remain the same, so $\tau_{\mathbb{C}_3}(R) = R$. Constructing τ with these clusters, the high level variables are revised to $\mathbf{V}_H = \{R, X, Y\}$, and $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$ can be computed as shown in the table in Fig. 8b. The C-DAG \mathcal{G}_C would be updated to the one in Fig. 8a.

It turns out that now, $Q = P(Y_{A=1, B=1} = 1)$ is τ -ID, and this can be shown by applying the backdoor-criterion (Pearl 2000, Thm. 3.3.2), adjusting over the variable R as follows.

$$\begin{aligned}
& P(Y_{A=1, B=1} = 1) \\
&= \sum_r P(Y = 1 \mid r, A = 1, B = 1)P(r) \\
&= \sum_r P(Y = 1 \mid r, X = 1)P(r). \tag{44}
\end{aligned}$$

Then $\sum_r P(Y = 1 \mid r, X = 1)P(r)$ could immediately be computed from $P(\mathbf{V}_H)$.

As an example, consider the following choice of \mathcal{M}_H , designed in a systematic way to match $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$.

$$\mathbf{U}_H = \{U_{RY}, U_{X0}, U_{X1}, U_{Y0}, U_{Y1}, U_{Y2}, U_{Y3}\} \tag{45}$$

$$\mathbf{V}_H = \{R, X, Y\} \tag{46}$$

$$\mathcal{F}_H = \begin{cases} R \leftarrow f_R(u_{RY}) = u_{RY} \\ X \leftarrow f_X(r, u_{X0}, u_{X1}) = \begin{cases} u_{X0} & r = 0 \\ u_{X1} & r = 1 \end{cases} \\ Y \leftarrow f_Y(x, u_{RY}, u_{Y0}, u_{Y1}, u_{Y2}, u_{Y3}) \\ = \begin{cases} u_{Y0} & x = 0, u_{RY} = 0 \\ u_{Y1} & x = 0, u_{RY} = 1 \\ u_{Y2} & x = 1, u_{RY} = 0 \\ u_{Y3} & x = 1, u_{RY} = 1 \end{cases} \end{cases} \tag{47}$$

$$P(\mathbf{U}_H) = \begin{cases} P(U_{RY} = 1) = 0.5 \\ P(U_{X0} = 1) = 0.04 \\ P(U_{X1} = 1) = 0.64 \\ P(U_{Y0} = 1) = 0.1 \\ P(U_{Y1} = 1) = 0.1 \\ P(U_{Y2} = 1) = 0.1 \\ P(U_{Y3} = 1) = 0.9 \end{cases} \tag{48}$$

One can verify that, indeed, $P^{\mathcal{M}_H}(\mathbf{V}_H)$ matches $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$ from Fig. 8b. It is also clear that \mathcal{M}_H induces \mathcal{G}_C . Further, we can compute the query to find that

$$\begin{aligned}
& P^{\mathcal{M}_H}(Y_{X=1} = 1) \\
&= P(U_{RY} = 0, U_{Y2} = 1) + P(U_{RY} = 1, U_{Y3} = 1) \\
&= (0.5)(0.1) + (0.5)(0.9) \\
&= 0.5, \tag{49}
\end{aligned}$$

which matches the true value $P^{\mathcal{M}_L}(Y_{A=1, B=1} = 1) = 0.5$ from Eq. 16. ■

The definition of τ -ID provides rigorous semantics to answer whether a query can be inferred across abstractions. The next step is to establish an approach to determine τ -ID when given the available data and graph. For this purpose, one fundamental result is that the notion of τ -ID is actually equivalent to classical identification in the higher level space, as shown by the following proposition.

Theorem 1 (Dual Abstract ID). *Consider a counterfactual query Q over \mathbf{V}_L , a constructive abstraction function τ w.r.t. clusters \mathbb{C} and \mathbb{D} , a C-DAG \mathcal{G}_C , and data \mathbb{Z} from \mathbf{V}_L . Q is τ -ID from \mathcal{G}_C and \mathbb{Z} if and only if $\tau(Q)$ is ID from \mathcal{G}_C and $\tau(\mathbb{Z})$. ■*

In words, τ -ID and classical ID on the high level space are equivalent. This is a powerful result since it implies that inferences can be made about the low level space by using existing results in the high level space. Our goal is to learn a higher level SCM \mathcal{M}_H to make inferences about \mathcal{M}_L , and we build on the machinery of Neural Causal Models (NCMs) (Xia et al. 2021) toward this goal. NCMs allow one to take the graph \mathcal{G}_C as an inductive bias (a \mathcal{G}_C -NCM as described in Def. 4) and leverage gradient-based methods to fit any SCM within the constrained space. Indeed, identification in NCMs can be shown to be equivalent to classical identification when considering models of the same granularity (Xia, Pan, and Bareinboim 2023, Thm. 3). When combined with Thm. 1, this implies the following result.

Corollary 1 (Abstract ID with NCMs). *Q is τ -ID from \mathcal{G}_C and \mathbb{Z} iff $\tau(Q)$ is Neural-ID from $\widehat{\Omega}(\mathcal{G}_C)$ and $\tau(\mathbb{Z})$. Moreover, if it is ID, then Q can be computed by $\tau(Q)$ from any \mathcal{G}_C -NCM \widehat{M} that is $\tau(\mathbb{Z})$ -consistent. ■*

In words, determining τ -ID is equivalent to determining neural identification (identification in the space of NCMs) on the higher level space. Further, to evaluate Q in identifiable cases, $\tau(Q)$ can be queried from any \mathcal{G}_C -NCM \widehat{M} that is $\tau(\mathbb{Z})$ -consistent. Corol. 1 implies that we can perform causal identification and estimation across abstractions using the NeuralID algorithm (Xia, Pan, and Bareinboim 2023, Alg. 1) on the high level space. This procedure is shown in Alg. 2. First, τ is constructed as described in Def. 6 given the clusters. Then, a \mathcal{G}_C -NCM is constructed over high-level variables \mathbf{V}_H . Two parameterizations of the NCM are created. Both are optimized to fit the transformed data $\tau(\mathbb{Z})$, but one is optimized to maximize the transformed query $\tau(Q)$ while the other is optimized to minimize it. If both parameterizations return the same result, then it must be the true value of the query; otherwise, the query is not identifiable.

To implement this algorithm in practice, we leverage the GAN-NCM approach introduced in Xia, Pan, and Bareinboim (2023); see details in Appendix C. Alg. 2 is sound and complete for solving the abstract identification problem, as shown below.

Corollary 2 (Soundness and Completeness). *Let \mathcal{M}_L be the low-level SCM, \mathbb{C} and \mathbb{D} be inter/intravariabile clusters of \mathbf{V}_L , \mathcal{G}_C be a C-DAG, Q be a query, and \widehat{Q} be the result from running Alg. 2 with inputs $\mathbb{Z}(\mathcal{M}_L) > 0$, \mathbb{C} , \mathbb{D} , \mathcal{G}_C , and Q . Then, Q is τ -ID from \mathcal{G}_C and \mathbb{Z} if and only if \widehat{Q} is not FAIL. Moreover, if \widehat{Q} is not FAIL, then $\widehat{Q} = Q(\mathcal{M}_L)$. ■*

While Alg. 2 solves the abstract ID problem, the consequences of the results in this section are more general. Notably, if Q is indeed τ -ID (which can be verified through Alg. 2), the algorithm produces a neural model \widehat{M} that serves as a proxy SCM that is Q - τ consistent with the true model \mathcal{M}_L . Such an SCM could serve as a generative model of the distribution Q , which has many uses. The samples generated from such a model could be used to estimate the query, or, in more complex settings such as with image data, it may be desirable to simply have novel generated samples consistent with the causal invariances embedded in the system.

Algorithm 2: NeuralAbstractID – Identifying and estimating queries across abstractions using NCMs.

Input : query Q , \mathcal{L}_2 datasets $\mathbb{Z}(\mathcal{M}_L)$, C-DAG \mathcal{G}_C , and admissible inter/intravariabile clusters \mathbb{C} and \mathbb{D} satisfying invariance condition

Output : $Q(\mathcal{M}_L)$ if identifiable, FAIL otherwise.

```

1  $\mathbf{V}_H \leftarrow \mathbb{C}, \mathcal{D}_{\mathbf{V}_H} \leftarrow \mathbb{D}$ 
2  $\tau \leftarrow \text{AbsFunc}(\mathbb{C}, \mathbb{D})$  // from Def. 6
3  $\widehat{M} \leftarrow \text{NCM}(\mathbf{V}_H, \mathcal{G}_C)$  // from Def. 4
4  $\theta_{\min}^* \leftarrow \arg \min_{\theta} \tau(Q)(\widehat{M}(\theta))$  s.t.
    $\tau(\mathbb{Z})(\widehat{M}(\theta)) = \tau(\mathbb{Z}(\mathcal{M}_L))$ 
5  $\theta_{\max}^* \leftarrow \arg \max_{\theta} \tau(Q)(\widehat{M}(\theta))$  s.t.
    $\tau(\mathbb{Z})(\widehat{M}(\theta)) = \tau(\mathbb{Z}(\mathcal{M}_L))$ 
6 if  $\tau(Q)(\widehat{M}(\theta_{\min}^*)) \neq \tau(Q)(\widehat{M}(\theta_{\max}^*))$  then
7 | return FAIL
8 else
9 | return  $\tau(Q)(\widehat{M}(\theta_{\min}^*))$  // choose min or
   |  $\tau(Q)(\widehat{M}(\theta_{\max}^*))$  // max arbitrarily

```

4 Representations in Learning Abstractions

In many applications, the choice of intervariable clusters \mathbb{C} is natural and can be made in tandem when deciding the assumptions of the C-DAG \mathcal{G}_C ¹¹. On the other hand, fully specifying the intravariabile clusters \mathbb{D} is usually challenging when working with high-dimensional data like images. If an intervariable cluster contained three binary variables from \mathbf{V}_L , then specifying its intravariabile cluster would require specifying some partition over its eight values, $(0, 0, 0), (0, 0, 1), \dots, (1, 1, 1)$, which is not difficult. However, if an intervariable cluster contained, for example, 128×128 pixels, each with 256 possible values, then the size of the domain of this cluster would be $256^{128 \times 128}$. Specifying an arbitrary partition over this many values is infeasible, as doing so would require an enumeration of every possible image along with some label designating each one to a cluster. In this section, we investigate the problem of learning abstractions when the intravariabile clusters \mathbb{D} are left unspecified.

While coarser clusters tend to be better in practice due to the dimensionality reduction, the theory in this paper can be applied for any choice of \mathbb{D} so long as the AIC (Def. 8) holds. Hence, a possible constraint when learning \mathbb{D} is to find a set of clusters such that the AIC is not violated. To this effect, the following result can be leveraged.

Proposition 5. *Consider a low level SCM \mathcal{M}_L and constructive abstraction function τ w.r.t. clusters \mathbb{C} and \mathbb{D} . \mathcal{M}_L is guaranteed to satisfy the AIC w.r.t. τ if and only if $\mathbb{D}_{\mathbf{C}_i} = \{\{\mathbf{c}_i\} : \mathbf{c}_i \in \mathcal{D}_{\mathbf{C}_i}\}$ for all $\mathbf{C}_i \in \mathbb{C}$. ■*

In words, the AIC is satisfied when the intravariabile clusters are maximal (i.e. each value is assigned to its own cluster). Consequently, this means that Alg. 2 can be applied in

¹¹Please refer to App. D.1 for best practices on how to choose or learn intervariable clusters when they are not given.

any case where τ_{C_i} is a bijective mapping between \mathcal{D}_{C_i} and $\mathcal{D}_{V_{H,i}}$. Also implied by this result is that, without additional information, one cannot choose any coarser clustering without potentially violating the AIC¹². For some intuition on why this is the case, consider the following example.

Example 9 (Example 1 continued). Consider once again the nutrition example where carbohydrates (C), fat (F), and protein (P) were combined to form the high-level variable, calories (Z). Calories was defined to be $Z = 4C + 9F + 4P$, resulting in a set of intravariabile clusters where (c_1, f_1, p_1) and (c_2, f_2, p_2) were clustered together if $4c_1 + 9f_1 + 4p_1 = 4c_2 + 9f_2 + 4p_2$.

It turns out that, without additional information, it is possible that this choice of clustering violates the AIC. For example, it could be the case that in the true model \mathcal{M}_L , the function for BMI f_B depends heavily on protein, since muscle density may affect BMI more than fat. In an extreme example, suppose

$$B \leftarrow f_B(c, f, p, u_B) = p + u_B, \quad (50)$$

that is, f_B only depends on protein out of the three macronutrients. In this case, $(C = 0, F = 0, P = 20)$ and $(C = 20, F = 0, P = 0)$ would result in different values of B despite being clustered together in the same intravariabile cluster, therefore, violating the AIC. Without information about f_B , it is not known whether any given two tuples, (c_1, f_1, p_1) and (c_2, f_2, p_2) , would witness this violation, so the only option is to leave all values of (C, F, P) in their own intravariabile clusters. ■

This example illustrates that without further information about the functions of the underlying SCM, it is, in general, impossible to cluster two intravariabile values together without potentially violating the AIC. This is also the case with the cholesterol discussion in Ex. 4 and 5. Without knowing f_Y , it would be impossible to determine whether the cluster choice in Ex. 5 would work better than the one in Ex. 4.

Still, Prop. 5 states that the intravariabile clustering which leaves every value in its own cluster will always satisfy the AIC. While this choice of \mathbb{D} does not reduce the dimensionality of the abstracted space, this means that we are not restricted to the original space of \mathbf{V}_L and can choose any \mathbf{V}_H with the same cardinality. In practice, this means that we can choose the option for \mathbf{V}_H that is the most beneficial for a given task. For example, some choices of \mathbf{V}_H and corresponding domains \mathcal{D}_{V_H} may have desirable properties such as simpler gradient computation, disentangled variables, or compatibility with arithmetic operations. In order to leverage this insight, we introduce the representational NCM.

Definition 11 (Representational NCM (RNCM)). A representational NCM (RNCM) is a tuple $\langle \hat{\tau}, \widehat{M} \rangle$, where $\hat{\tau}(\mathbf{v}_L; \theta_\tau)$ is a function parameterized by θ_τ mapping from \mathbf{V}_L to \mathbf{V}_H , and \widehat{M} is an NCM defined over \mathbf{V}_H . A \mathcal{G}_C -constrained

¹²In many cases, there may be additional information in the form of invariances (e.g. rotational invariance in image data). In such cases, this information can be leveraged to learn coarser clusters. See Appendix D.3 for more details.

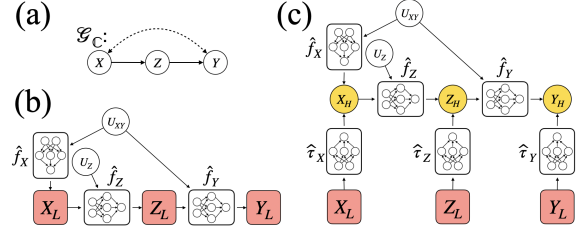


Figure 9: Example using graph \mathcal{G}_C shown in (a) to compare (b) the \mathcal{G}_C -NCM and (c) \mathcal{G}_C -RNCM. Functions of the NCM directly output values of the lower level variables (grouped by clusters in \mathbb{C}), while functions of the RNCM output values of their higher level counterparts, mapped by $\hat{\tau}$.

RNCM (\mathcal{G}_C -RNCM) is an RNCM $\langle \hat{\tau}, \widehat{M} \rangle$ such that $\hat{\tau}$ is composed of subfunctions $\hat{\tau}_{C_i}$ for each $C_i \in \mathbb{C}$ (each with its own parameters $\theta_{\tau_{C_i}}$), and \widehat{M} is a \mathcal{G}_C -NCM (Def. 4). ■

In words, an RNCM is a pair of a parameterized abstraction function $\hat{\tau}$ and an NCM \widehat{M} defined over the space of high level variables \mathbf{V}_H obtained from $\hat{\tau}(\mathbf{V}_L)$. The \mathcal{G}_C -RNCM is simply an RNCM constrained over the C-DAG \mathcal{G}_C , where $\hat{\tau}$ must use the intervariable clusters specified by \mathbb{C} . Fig. 9 illustrates the difference between a \mathcal{G}_C -RNCM and a standard \mathcal{G}_C -NCM. Given \mathcal{G}_C in (a), the \mathcal{G}_C -NCM (b) directly defines the domains of the inputs and outputs of the functions to be the space of the variables \mathbf{V}_L . On the other hand, the \mathcal{G}_C -RNCM (c) defines the domains over $\mathbf{V}_H = \hat{\tau}(\mathbf{V}_L)$ instead.

Training can be done as a two-step process, where first $\hat{\tau}$ is trained to map to an optimal task-specific space, and then \widehat{M} can be trained on $\hat{\tau}(\mathbf{V}_L)$ (e.g., through Alg. 2). In the first step, $\hat{\tau}$ (parameterized by θ_τ) can be trained according to a loss such as

$$L_{\hat{\tau}}(\mathbf{v}_L) = \|\hat{\tau}^{-1}(\hat{\tau}(\mathbf{v}_L, \theta_\tau); \theta_{\tau-1}) - \mathbf{v}_L\|^2 + \lambda_r L_r(\mathbf{v}_L), \quad (51)$$

where $\hat{\tau}^{-1}$ is a neural network parameterized by $\theta_{\tau-1}$ that attempts to invert $\hat{\tau}$ and recover the original \mathbf{v}_L , L_r is a separate representation regularizer, and λ_r is its regularization strength. The first term in Eq. 51 is used to enforce bijectivity between \mathcal{D}_{C_i} and $\mathcal{D}_{V_{H,i}}$, as required by Prop. 5. In words, one can train $\hat{\tau}$ in an autoencoder-like setup (Kramer 1991; Kingma and Welling 2014) with a reconstruction loss. The second term is left open-ended and can be used to impose a desired form on the output of $\hat{\tau}$. The following example helps to ground this point.

Example 10. Consider a simple example with $\mathbf{V}_L = \mathbf{X} \cup \{Y\}$, where \mathbf{X} is the collection of pixels representing an image of a cat or dog, and Y is a binary label (caused by \mathbf{X}) predicting whether the animal depicted in \mathbf{X} is aggressive or docile. Suppose researchers are studying the relationship between \mathbf{X} and Y and would like to work in a more abstract space. Denote X_H as the high-level counterpart of \mathbf{X} . Moreover, suppose the researchers are given another set of labels of \mathbf{X} , say Z , which state whether the animal in the image is a cat or dog. Z is not included in the study with \mathbf{V}_L , but it could potentially be used for learning the space of X_H .

Specifically, the researchers would like to train $\hat{\tau}_X : \mathcal{D}_X \rightarrow \mathcal{D}_{X_H}$ such that the AIC is satisfied from Prop. 5, and X_H additionally encodes some information about Z . One clever approach is to introduce another neural network $g(x_H; \theta_g)$ parameterized by θ_g that aims to classify Z from X_H . The researchers could train $\hat{\tau}_X$ with the loss $L_{\hat{\tau}}$ from Eq. 51 such that L_r is the classification loss of g . This would result in X_H encoding both enough information to reconstruct \mathbf{X} and also to classify Z . Although X_H may not have lower dimensionality than the original X , it may be more useful in a downstream task, such as building a classifier for Y (i.e. X_H may be a more well-behaved set of features for a classifier of Y than the original set of pixel values \mathbf{X}). ■

In general, $\hat{\tau}$ can be thought of as a function mapping to a representation space, where the second term in Eq. 51 can be used to regularize the representation space for a desired task. The flexibility of this approach makes it amenable to the wide developments of the representation learning literature (Bengio, Courville, and Vincent 2013). We empirically demonstrate this approach below in the experiment of Sec. 5.2.

5 Experiments

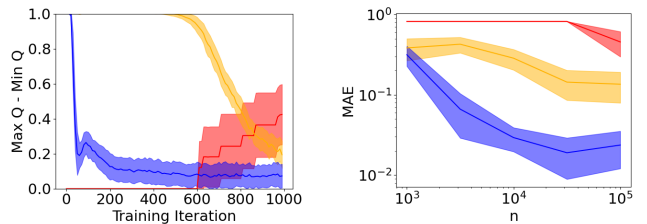
In this section, we empirically evaluate the effects of utilizing abstractions in causal inference tasks. More details on the data-generating models and architectures can be found in Appendix C. Implementation code is publicly available at <https://github.com/CausalAILab/NeuralCausalAbstractions>.

5.1 Nutritional Study

We follow up on the nutrition study discussed in Ex. 1. Since a BMI of 25 or over is considered overweight, the goal is to identify and estimate the query $Q = P(B_{D=d} \geq 25)$, the causal effect of diet on weight, given the available graphical constraints and observational data $P(\mathbf{V}_L)$ using Alg. 2. R and D are 32-dimensional one-hot vectors, and the others are real-valued, so the query may be difficult to answer given such high-dimensional variables. Instead, it may be more effective to work in an abstract space with the proposed inter-variable clusters $\mathbb{C} = \{D_H = \{D\}, Z = \{C, F, P\}, B_H = \{B\}\}$. The original graph \mathcal{G} and corresponding C-DAG \mathcal{G}_C are shown in Fig. 5. We are also given intravariabe clusters \mathbb{D} such that all values of D_H, Z , and B_H are clustered into binary categories. Specifically, $D_H = 1$ denotes unhealthy dishes, $Z = 1$ denotes high calorie count, and $B_H = 1$ denotes an overweight BMI (≥ 25).

We compare the effectiveness identifying and estimating Q with NCMs in three different settings, with results shown in Fig. 10. The first approach (red) attempts to solve the problem directly in the space of \mathbf{V}_L by identifying and estimating Q from the original causal diagram \mathcal{G} and observational dataset from $P(\mathbf{V}_L)$. The second approach (yellow) solves the same task but first normalizes each variable¹³ of the data between 0 and 1. The third approach (blue) is the newly

¹³This normalization approach is equivalent to using a constructive abstraction function τ over the full set of clusters $\mathbb{C} = \mathbf{V}_L$ and $\mathbb{D} = \mathcal{D}_{\mathbf{V}_L}$ (with a bit of abuse of notation). That is, each variable and value are placed in their own cluster, and all values are simply remapped to different values.



(a) Gaps between max and min query across 1000 training iterations when running Alg. 2. (b) Mean absolute error (MAE) v. dataset size (in log-log scale) for query estimation.

Figure 10: Results of the nutrition experiment. Our approach (blue) is compared with a GAN-NCM trained on raw data (red) and one trained on normalized data (yellow).




Figure 11: Colored MNIST Experimental Setup

proposed approach and leverages the concept of τ -ID, identifying and estimates Q from the C-DAG \mathcal{G}_C and high level data $\tau(P(\mathbf{V}_L))$. The model is trained over the abstract space of \mathbf{V}_H computed using the constructive abstraction function τ defined on \mathbb{C} and \mathbb{D} . All three approaches are implemented in the style of GAN-NCM (Xia, Pan, and Bareinboim 2023). Since Q is identifiable, the gap between the max and min queries computed in Alg. 2 are expected to be as small as possible. As shown in Fig. 10a, the proposed approach (blue) converges quickly while others fail to close the gap between the max and min queries. Fig. 10b also shows that the proposed approach can estimate Q with significantly lower error. Furthermore, since the proposed approach uses the C-DAG \mathcal{G}_C instead of the original causal diagram \mathcal{G} , the approach operates under fewer assumptions of domain knowledge.

5.2 Colored MNIST Digits

We evaluate the RNCM in a high-dimensional image dataset of colorized MNIST (Deng 2012) digits. Each image (I) has a corresponding digit (D) and color (C) label, and their relationships are shown in the C-DAG \mathcal{G}_C in Fig. 11a. Color and digit are highly correlated (e.g. 0s are typically red, while 5s are cyan), as shown in Fig. 11b. We evaluate three approaches in the task of sampling images from causal queries. The first approach is a naïve conditional GAN that does not take causality into account. The second is a standard GAN-NCM as described in Xia, Pan, and Bareinboim (2023). The third is called a GAN-RNCM, a GAN implementation of the representational NCM following the approach described in Sec. 4.

Samples of the results are shown in Fig. 12. All models are capable of producing digit images, as shown in the first column. The second column illustrates $P(I | D = 0)$, the images conditioned on digit = 0. Many red 0s are expected

	Image Samples	$P(I D = 0)$	$P(I_{D=0})$	$P(I_{D=0} D = 5)$
Conditional GAN				
GAN-NCM				
GAN-RNCM (Ours)				
Ground Truth				

Figure 12: Colored MNIST results. Samples from various causal queries (top) are collected from competing approaches (left), with the ground truth samples from the data generating model shown in the bottom row. The left column simply shows image samples from $P(I)$ from each of the models, while the second, third, and fourth columns show samples generated from an \mathcal{L}_1 , \mathcal{L}_2 , and \mathcal{L}_3 query, respectively.

since most 0s are red in the dataset. The third column illustrates the interventional query $P(I_{D=0})$, the images with digits forced to be 0 through intervention. As interventions ignore the spurious correlations between color and digit, 0s of all colors are expected. Finally, the fourth column illustrates the counterfactual query $P(I_{D=0} | D = 5)$, indicating what the digits would have looked like had they been 0, given that they were originally 5. Since 5s tend to be cyan, the samples are expected to be 0s that retain the cyan color of the 5s.

In all cases, GAN-RNCM (new approach) produces results closer to the expected outcomes, as shown by the ground truth. The conditional GAN fails to distinguish causal queries from conditional queries, and samples appear similar in the 2nd, 3rd, and 4th columns. The standard GAN-NCM faces challenges disentangling color from digit, as shown from the presence of several non-zero digits in the 3rd column and several digits that resemble 5s in the 4th column. Further, both the conditional GAN and the GAN-NCM face challenges in capturing the relationship between color and image in more complex distributions, as evident from the mosaic coloring in many of the samples. The GAN-RNCM significantly outperforms the other approaches in terms of sample quality for the causal queries.

6 Conclusions

In this paper, we developed a new framework of abstractions based on the PCH’s layers with the goal of learning a high-level causal model at a coarser granularity. In each section of the paper, we relaxed certain assumptions that are not easily achievable in practice, and showed how to obtain the high-level model under these particular settings. These increasingly more refined results are summarized in Fig. 13.

We started by noting that previous works on causal abstractions, such as Beckers and Halpern (2019), set a strong foundation of defining abstractions in rigorous terms. (A detailed discussion is provided in Appendix B.) The main drawback is that these definitions are declarative, meaning that given a low level SCM \mathcal{M}_L , high level SCM \mathcal{M}_H , and an abstraction function τ , the definitions in these works can be used to decide whether \mathcal{M}_H is an abstraction of \mathcal{M}_L or not. Still, solving this decision task requires a substantial amount of input information, as highlighted in the first row of

Fig. 13, and which is unavailable in many practical settings.

We then examined in Sec. 2 the challenge of obtaining the abstract model \mathcal{M}_H when the same is unavailable, as shown in the second row of Fig. 13. We provided a systematic way of constructing τ given inter/intravariabile clusters (Defs. 5 and 6). Building on this, we developed Alg. 1, which allows one to obtain the higher level abstraction \mathcal{M}_H given \mathcal{M}_L and the constructive abstraction function τ .

Still, these results can be hard to apply since \mathcal{M}_L is not commonly available in some real world settings. Rather, partial information about \mathcal{M}_L through its data distributions may be available. We then described in Sec. 3 how to obtain \mathcal{M}_H under these restrictions, as shown in the third row of Fig. 13. Causal inferences on higher levels of the PCH cannot be realized using lower layers alone, as shown by Prop. 4, which means that assumptions are needed. The assumptions considered in this paper take the form of C-DAGs (Def. 9), an abstract version of causal diagrams leveraging the intervariable clusters. Then, given the data \mathbb{Z} , the C-DAG \mathcal{G}_C , and the abstraction function τ constructed from the clusters, Alg. 2 can be used to learn the abstract NCM $\widehat{\mathcal{M}}_H$. Then, $\widehat{\mathcal{M}}_H$ can be used for tasks such as identification, estimation, or sampling.

We showed in Sec. 4 how to learn τ , noting that acquiring intravariabile clusters \mathbb{D} may be challenging in practice. The fourth row of Fig. 13 highlight this task. Specifically, we introduced the representational NCM or RNCM (Def. 11), which parameterizes $\widehat{\tau}$ as a neural network. We then developed a learning procedure for intravariabile clusters, considering task-specific objectives. Whenever additional information about the problem is available in the form of invariances (e.g., translational, rotational, permutation), coarser clusters can be learned, as elaborated in Appendix D.3.

Finally, although not shown in the table, additional technical content can be found in Appendix D. Specifically, Appendix D.1 discusses how to make the best choice of intervariable clusters \mathbb{C} when it is not given. Appendix D.2 describes possible ways to relax the AIC (Def. 8), which is an assumption that is made throughout the paper. We encourage further research on the topics covered in this paper, such as on the best way to learn $\widehat{\tau}$ in an RNCM.

Approach	SCM \mathcal{M}_L				Abstraction τ			SCM \mathcal{M}_H	Output
	\mathcal{M}_L	Data \mathbb{Z}	Graph \mathcal{G}	C-DAG \mathcal{G}_C	τ	\mathbb{C}	\mathbb{D}	\mathcal{M}_H	
Existing works	✓	✓	✓	✓	✓	✓	✓	✓	Abstraction: Yes/No
Sec. 2	✓	✓	✓	✓	-	✓	✓	-	\mathcal{L}_3 - τ consistent \mathcal{M}_H
Sec. 3	-	✓	-	✓	-	✓	✓	-	\mathbb{Z} - τ / \mathcal{G}_C consistent \mathcal{M}_H
Sec. 4	-	✓	-	✓	-	✓	-	-	\mathbb{Z} - τ / \mathcal{G}_C consistent \mathcal{M}_H

Figure 13: Summary table of contributions of each section in terms of input assumptions and outputs. A checkmark (✓) indicates that the corresponding information is assumed to be available in that section. A gray checkmark indicates that it is implied by a stronger assumption (e.g. data \mathbb{Z} can be sampled from \mathcal{M}_L if \mathcal{M}_L is available). A dash (-) indicates that the information is not assumed. An approach that has fewer checkmarks makes fewer assumptions and is therefore more applicable in practice.

Acknowledgements

This research was supported in part by the NSF, ONR, AFOSR, DARPA, DoE, Amazon, JP Morgan, and The Alfred P. Sloan Foundation.

References

- Anand, T. V.; Ribeiro, A. H.; Tian, J.; and Bareinboim, E. 2023. Causal Effect Identification in Cluster DAGs. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*. AAAI Press.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*, 214–223. PMLR.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Bareinboim, E.; Correa, J. D.; Ibeling, D.; and Icard, T. 2022. On Pearl’s Hierarchy and the Foundations of Causal Inference. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, 507–556. New York, NY, USA: Association for Computing Machinery, 1st edition.
- Bareinboim, E.; and Pearl, J. 2016. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113(27): 7345–7352.
- Beckers, S.; Eberhardt, F.; and Halpern, J. Y. 2019. Approximate Causal Abstraction. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence*.
- Beckers, S.; and Halpern, J. Y. 2019. Abstracting Causal Models. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI’19/IAAI’19/EAAI’19. AAAI Press. ISBN 978-1-57735-809-1.
- Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation Learning: A Review and New Perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(8): 1798–1828.
- Brock, A.; Donahue, J.; and Simonyan, K. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations*.
- Chalupka, K.; Perona, P.; and Eberhardt, F. 2015. Visual Causal Feature Learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, UAI’15, 181–190. Arlington, Virginia, USA: AUAI Press. ISBN 9780996643108.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. E. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 1597–1607. PMLR.
- Correa, J.; Lee, S.; and Bareinboim, E. 2021. Nested Counterfactual Identification from Arbitrary Surrogate Experiments. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 6856–6867. Curran Associates, Inc.
- Deng, L. 2012. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6): 141–142.
- Du, X.; Sun, L.; Duivesteyn, W.; Nikolaev, A.; and Pechenizkiy, M. 2020. Adversarial Balancing-based Representation Learning for Causal Effect Inference with Observational Data. arXiv:1904.13335.
- Falcon, W.; and Cho, K. 2020. A Framework For Contrastive Self-Supervised Learning And Designing A New Approach. *arXiv preprint arXiv:2009.00104*.
- Felekis, Y.; Zennaro, F. M.; Branchini, N.; and Damoulas, T. 2024. Causal Optimal Transport of Abstractions. In *Conference on Causal Learning and Reasoning, CLear 2024*.
- Galles, D.; and Pearl, J. 1998. An axiomatic characterization of causal counterfactuals. *Foundation of Science*, 3(1): 151–182.
- Gamba, R.; Schuchter, J.; Rutt, C.; and Seto, E. 2014. Measuring the Food Environment and its Effects on Obesity in the United States: A Systematic Review of Methods and Results. *Journal of Community Health*, 40(3): 464–475.
- Geiger, A.; Potts, C.; and Icard, T. 2023. Causal Abstraction for Faithful Model Interpretation. arXiv:2301.04709.
- Glorot, X.; and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In Teh, Y. W.; and Titterton, M., eds., *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, 249–256. Chia Laguna Resort, Sardinia, Italy: PMLR.

- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems*, volume 27, 2672–2680. Curran Associates, Inc.
- Goudet, O.; Kalainathan, D.; Caillou, P.; Guyon, I.; Lopez-Paz, D.; and Sebag, M. 2018. Learning functional causal models with generative neural networks. In *Explainable and interpretable models in computer vision and machine learning*, 39–80. Springer.
- Graves, A.; and Jaitly, N. 2014. Towards End-To-End Speech Recognition with Recurrent Neural Networks. In Xing, E. P.; and Jebara, T., eds., *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, 1764–1772. Beijing, China: PMLR.
- Guo, R.; Cheng, L.; Li, J.; Hahn, P. R.; and Liu, H. 2020. A Survey of Learning Causality with Data. *ACM Computing Surveys*, 53(4): 1–37.
- Halpern, J. Y. 1998. Axiomatizing Causal Reasoning. In G.F. Cooper; and S. Moral, eds., *Uncertainty in Artificial Intelligence*, 202–210. Cornell University, San Francisco, CA: Morgan Kaufmann.
- Hinton, G.; Srivastava, N.; and Swersky, K. 2018. *Neural Networks for Machine Learning Lecture 6a Overview of mini-batch gradient descent*.
- Ibeling, D.; and Icard, T. 2020. Probabilistic reasoning across the causal hierarchy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10170–10177.
- Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 448–456. Lille, France: PMLR.
- Johansson, F. D.; Shalit, U.; and Sontag, D. 2016. Learning Representations for Counterfactual Inference. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, 3020–3029. JMLR.org.
- Jung, Y.; Tian, J.; and Bareinboim, E. 2020a. Estimating Causal Effects Using Weighting-Based Estimators. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*. AAAI Press.
- Jung, Y.; Tian, J.; and Bareinboim, E. 2020b. Learning Causal Effects via Weighted Empirical Risk Minimization. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 12697–12709. Vancouver, Canada: Curran Associates, Inc.
- Jung, Y.; Tian, J.; and Bareinboim, E. 2021. Estimating Identifiable Causal Effects through Double Machine Learning. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, R-69. Vancouver, Canada: AAAI Press.
- Kallus, N. 2020. DeepMatch: Balancing Deep Covariate Representations for Causal Inference Using Adversarial Training. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 5067–5077. PMLR.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In Bengio, Y.; and LeCun, Y., eds., *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Kocaoglu, M.; Snyder, C.; Dimakis, A. G.; and Vishwanath, S. 2018. CausalGAN: Learning Causal Implicit Generative Models with Adversarial Training. In *International Conference on Learning Representations*.
- Kramer, M. A. 1991. Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, 37(2): 233–243.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In Pereira, F.; Burges, C. J. C.; Bottou, L.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems*, volume 25, 1097–1105. Curran Associates, Inc.
- LeCun, Y.; and Bengio, Y. 1998. *Convolutional Networks for Images, Speech, and Time Series*, 255–258. Cambridge, MA, USA: MIT Press. ISBN 0262511029.
- Li, S.; and Fu, Y. 2017. Matching on Balanced Nonlinear Representations for Treatment Effects Estimation. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30, 929–939. Curran Associates, Inc.
- Louizos, C.; Shalit, U.; Mooij, J.; Sontag, D.; Zemel, R.; and Welling, M. 2017. Causal Effect Inference with Deep Latent-Variable Models. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, 6449–6459. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510860964.
- Massidda, R.; Geiger, A.; Icard, T.; and Bacciu, D. 2023. Causal Abstraction with Soft Interventions. In van der Schaar, M.; Zhang, C.; and Janzing, D., eds., *Proceedings of the Second Conference on Causal Learning and Reasoning*, volume 213 of *Proceedings of Machine Learning Research*, 68–87. PMLR.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing Atari With Deep Reinforcement Learning. In *NIPS Deep Learning Workshop*.
- Murphy, R. L.; Srinivasan, B.; Rao, V.; and Ribeiro, B. 2019. Janossy Pooling: Learning Deep Permutation-Invariant Functions for Variable-Size Inputs. In *International Conference on Learning Representations*.

- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch.
- Pearl, J. 1995. Causal diagrams for empirical research. *Biometrika*, 82(4): 669–688.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. New York, NY, USA: Cambridge University Press, 2nd edition.
- Pearl, J.; and Mackenzie, D. 2018. *The Book of Why*. New York: Basic Books.
- Rubenstein, P. K.; Weichwald, S.; Bongers, S.; Mooij, J.; Janzing, D.; Grosse-Wentrup, M.; and Schölkopf, B. 2017. Causal Consistency of Structural Equation Models. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence*.
- Shalit, U.; Johansson, F. D.; and Sontag, D. 2017. Estimating individual treatment effect: generalization bounds and algorithms. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 3076–3085. International Convention Centre, Sydney, Australia: PMLR.
- Shi, C.; Blei, D. M.; and Veitch, V. 2019. Adapting Neural Networks for the Estimation of Treatment Effects. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 2503–2513.
- Shorten, C.; and Khoshgoftaar, T. M. 2019. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data*, 6: 1–48.
- Spirtes, P.; and Scheines, R. 2004. Causal Inference of Ambiguous Manipulations. *Philosophy of Science*, 71: 833–845.
- Steinberg, D. 2007. Copyright. In *The Cholesterol Wars*. Oxford: Academic Press. ISBN 978-0-12-373979-7.
- Truswell, A. 2010. *Cholesterol and Beyond: The Research on Diet and Coronary Heart Disease 1900-2000*. ISBN 978-90-481-8874-1.
- Villani, C. 2009. *Optimal transport: old and new*, volume 338. Springer.
- Winston, M. E.; Chaffin, R.; and Herrmann, D. 1987. A Taxonomy of Part-Whole Relations. *Cognitive Science*, 11(4): 417–444.
- Xia, K.; Lee, K.-Z.; Bengio, Y.; and Bareinboim, E. 2021. The Causal-Neural Connection: Expressiveness, Learnability, and Inference. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 10823–10836. Curran Associates, Inc.
- Xia, K.; Pan, Y.; and Bareinboim, E. 2023. Neural Causal Models for Counterfactual Identification and Estimation. In *Proceedings of the 11th International Conference on Learning Representations (ICLR-23)*.
- Yao, L.; Li, S.; Li, Y.; Huai, M.; Gao, J.; and Zhang, A. 2018. Representation Learning for Treatment Effect Estimation from Observational Data. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31, 2633–2643. Curran Associates, Inc.
- Yoon, J.; Jordon, J.; and van der Schaar, M. 2018. GAN-ITE: Estimation of Individualized Treatment Effects using Generative Adversarial Nets. In *International Conference on Learning Representations*.
- Zaheer, M.; Kottur, S.; Ravanbakhsh, S.; Póczos, B.; Salakhutdinov, R.; and Smola, A. J. 2017. Deep Sets. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, 3394–3404. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510860964.
- Zennaro, F. M.; Drávcz, M.; Apachitei, G.; Widanage, W. D.; and Damoulas, T. 2023. Jointly Learning Consistent Causal Abstractions Over Multiple Interventional Distributions. In van der Schaar, M.; Zhang, C.; and Janzing, D., eds., *Conference on Causal Learning and Reasoning, CLeaR 2023, 11-14 April 2023, Amazon Development Center, Tübingen, Germany, April 11-14, 2023*, volume 213 of *Proceedings of Machine Learning Research*, 88–121. PMLR.

A Proofs

A.1 Counterfactual Axioms and Properties

The proofs in this work will rely on the following counterfactual axioms from Galles and Pearl (1998) and Halpern (1998):

Fact 1 ((Galles and Pearl 1998, Thms. 1,2)). *The following properties hold in SCMs:*

- 1 (Composition). For any two singleton variables Y and W , and any set of variables \mathbf{X} in a causal model, we have

$$W_{\mathbf{x}}(\mathbf{u}) = w \implies Y_{\mathbf{x},w}(\mathbf{u}) = Y_{\mathbf{x}}(\mathbf{u})$$

- 2 (Effectiveness) For all variables \mathbf{X} and \mathbf{W} , $\mathbf{X}_{\mathbf{x}\mathbf{w}}(\mathbf{u}) = \mathbf{x}$. ■

As a consequence of these axioms, we can prove the following result:

Corollary 3. *For any set of variables \mathbf{Y} and \mathbf{X} , we have*

$$\mathbf{Y}_{\mathbf{x}}(\mathbf{u}) = \mathbf{Y}_{\mathbf{pa}_{\mathbf{Y}}^{(1)}, \mathbf{pa}_{\mathbf{Y}[\mathbf{x}]}^{(2)}(\mathbf{u})}(\mathbf{u}) = \mathbf{Y}_{\mathbf{pa}_{\mathbf{Y}[\mathbf{x}]}(\mathbf{u})}(\mathbf{u}),$$

where $\mathbf{pa}_{\mathbf{Y}}^{(1)} = \mathbf{pa}_{\mathbf{Y}} \cap \mathbf{X}$, $\mathbf{pa}_{\mathbf{Y}}^{(2)}$ are its corresponding values of \mathbf{x} , and $\mathbf{pa}_{\mathbf{Y}}^{(2)} = \mathbf{pa}_{\mathbf{Y}} \setminus \mathbf{pa}_{\mathbf{Y}}^{(1)}$. ■

Proof. We have

$$\mathbf{Y}_{\mathbf{x}}(\mathbf{u}) = \mathbf{Y}_{\mathbf{pa}_{\mathbf{Y}}^{(1)}, \mathbf{x}'}(\mathbf{u}) \quad (52)$$

$$= \mathbf{Y}_{\mathbf{pa}_{\mathbf{Y}}^{(1)}, \mathbf{pa}_{\mathbf{Y}[\mathbf{x}]}^{(2)}(\mathbf{u}), \mathbf{x}'}(\mathbf{u}) \quad (53)$$

$$= \mathbf{Y}_{\mathbf{pa}_{\mathbf{Y}}^{(1)}, \mathbf{pa}_{\mathbf{Y}[\mathbf{x}]}^{(2)}(\mathbf{u})}(\mathbf{u}) \quad (54)$$

$$= \mathbf{Y}_{\mathbf{pa}_{\mathbf{Y}[\mathbf{x}]}^{(1)}(\mathbf{u}), \mathbf{pa}_{\mathbf{Y}[\mathbf{x}]}^{(2)}(\mathbf{u})}(\mathbf{u}) \quad (55)$$

$$= \mathbf{Y}_{\mathbf{pa}_{\mathbf{Y}[\mathbf{x}]}(\mathbf{u})}(\mathbf{u}). \quad (56)$$

Here, \mathbf{x}' denotes the values of \mathbf{x} corresponding to $\mathbf{X} \setminus \mathbf{pa}_{\mathbf{Y}}^{(1)}$. Eq. 52 holds by definition. Eq. 53 holds by the composition property, which can be applied iteratively to each variable in \mathbf{Y} and $\mathbf{pa}_{\mathbf{Y}}^{(2)}$. Eq. 54 holds because the value of \mathbf{Y} is deterministic once $\mathbf{pa}_{\mathbf{Y}}$ and \mathbf{U} , the inputs to its functions, are fixed. Eq. 55 holds by the effectiveness property. Finally Eq. 56 holds by definition. □

We also leverage the following results from Correa, Lee, and Bareinboim (2021):

Definition 12 ((Correa, Lee, and Bareinboim 2021, Def. 3)). The set of (counterfactual) ancestors of $Y_{\mathbf{x}}$ w.r.t. graph \mathcal{G} , denoted $An(Y_{\mathbf{x}})$, consist of each $W_{\mathbf{z}}$ such that $W \in An(Y)_{\mathcal{G}_{\mathbf{x}}}$ (which includes Y itself), and $\mathbf{z} = \mathbf{x} \cap An(W)_{\mathcal{G}_{\mathbf{x}}}$. ■

Fact 2 ((Correa, Lee, and Bareinboim 2021, Thm. 1)). *Let \mathbf{W}_* be an ancestral set, that is, $An(\mathbf{W}_*) = \mathbf{W}_*$, and let \mathbf{w}_* be a vector with a value for each variable in \mathbf{W}_* . Then,*

$$P(\mathbf{W}_* = \mathbf{w}_*) = P\left(\bigwedge_{W_t \in \mathbf{W}_*} W_{\mathbf{pa}_W} = w\right),$$

where each w is w_t and \mathbf{pa}_W is determined for each $W_t \in \mathbf{W}_*$ as follows:

- (i) the values for variables in $\mathbf{pa}_W \cap \mathbf{T}$ are the same as in \mathbf{t} , and
- (ii) the values for variables in $\mathbf{pa}_W \setminus \mathbf{T}$ are taken from \mathbf{w}_* , corresponding to the parents of W . ■

For the proofs in this work, we leverage a key concept of counterfactuals which we define as functional counterfactuals.

Definition 13 (Functional Counterfactuals). For any SCM $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$, denote

$$\mathbf{F} = \left\{ V_{i[\mathbf{pa}_{V_i}^{(j)}]} : V_i \in \mathbf{V}, \mathbf{pa}_{V_i}^{(j)} \in \mathcal{D}_{\mathbf{pa}_{V_i}} \right\} \quad (57)$$

as the functional counterfactual set of \mathcal{M} , a set of counterfactual variables containing each variable intervened on every possible instantiation of its parents. Denote \mathbf{f} and $\mathcal{D}_{\mathbf{F}}$ as its instantiation and domain respectively¹⁴. If $\mathbf{pa}_{V_i} = \emptyset$, then $V_i \in \mathbf{F}$ with no intervention. ■

The idea behind this definition is to establish a standard family of counterfactual quantities which generalizes all other counterfactuals. We will see the power of functional counterfactuals through the following lemmas.

Lemma 1 (Functional Counterfactual Uniqueness). *Let $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$ be an SCM with functional counterfactual set \mathbf{F} . Let $\mathcal{D}_{\mathbf{U}}(\mathbf{f}) \subseteq \mathcal{D}_{\mathbf{U}}$ be the set of values of \mathbf{U} such that for every $\mathbf{u} \in \mathcal{D}_{\mathbf{U}}(\mathbf{f})$, we have $\mathbf{F} = \mathbf{f}$ when evaluating \mathcal{M} with $\mathbf{U} = \mathbf{u}$. Then, for any $\mathbf{u} \in \mathcal{D}_{\mathbf{U}}$, there exists a unique $\mathbf{f} \in \mathcal{D}_{\mathbf{F}}$ such that $\mathbf{u} \in \mathcal{D}_{\mathbf{U}}(\mathbf{f})$. ■*

Proof. Note that for any $\mathbf{u} \in \mathcal{D}_{\mathbf{U}}$, we can construct the value of \mathbf{f} such that $\mathbf{u} \in \mathcal{D}_{\mathbf{U}}(\mathbf{f})$ as follows. For every $V_i \in \mathbf{V}$ and $\mathbf{pa}_{V_i}^{(j)} \in \mathcal{D}_{\mathbf{pa}_{V_i}}$, choose $v_i^{(j)} = f_{V_i}(\mathbf{pa}_{V_i}^{(j)}, \mathbf{u})$. Collectively, these values of $v_i^{(j)}$ can be used to form \mathbf{f} . It is clear that $\mathbf{u} \in \mathcal{D}_{\mathbf{U}}(\mathbf{f})$ because $V_{i[\mathbf{pa}_{V_i}^{(j)}]}(\mathbf{u}) = f_{V_i}(\mathbf{pa}_{V_i}^{(j)}, \mathbf{u}) = v_i^{(j)}$ for all i and j , implying that $\mathbf{F} = \mathbf{f}$. Further, \mathbf{f} is unique, since $V_{i[\mathbf{pa}_{V_i}^{(j)}]}(\mathbf{u})$ can only be equal to one unique deterministic value once \mathbf{u} is fixed. □

Lemma 2 (Functional Counterfactual Completeness). *Let $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$, $\mathcal{M}' = \langle \mathbf{U}', \mathbf{V}, \mathcal{F}', P(\mathbf{U}') \rangle$ be two SCMs both defined over \mathbf{V} and with the same functional counterfactual set \mathbf{F} . Then $P^{\mathcal{M}}(\mathbf{F} = \mathbf{f}) = P^{\mathcal{M}'}(\mathbf{F} = \mathbf{f})$ for all $\mathbf{f} \in \mathcal{D}_{\mathbf{F}}$ if and only if $\mathcal{L}_3(\mathcal{M}) = \mathcal{L}_3(\mathcal{M}')$. ■*

Proof. The backward direction of this proof is trivial since all functional counterfactuals belong to the set of all counterfactuals (i.e. $P^{\mathcal{M}}(\mathbf{F} = \mathbf{f}) \in \mathcal{L}_3(\mathcal{M})$ and $P^{\mathcal{M}'}(\mathbf{F} = \mathbf{f}) \in \mathcal{L}_3(\mathcal{M}')$ for all $\mathbf{f} \in \mathcal{D}_{\mathbf{F}}$). Therefore, $\mathcal{L}_3(\mathcal{M}) = \mathcal{L}_3(\mathcal{M}')$ implies $P^{\mathcal{M}}(\mathbf{F} = \mathbf{f}) = P^{\mathcal{M}'}(\mathbf{F} = \mathbf{f})$ for all $\mathbf{f} \in \mathcal{D}_{\mathbf{F}}$.

¹⁴We intentionally use the notation of \mathbf{F} because this set of counterfactual quantities is heavily related to the functions of the SCM, \mathcal{F} . Notably, the behavior of a function $f_{V_i} \in \mathcal{F}$ w.r.t. \mathbf{U} can be specified fully by the joint counterfactual $P\left(\bigwedge_{\mathbf{pa}_{V_i}^{(j)}} V_{i[\mathbf{pa}_{V_i}^{(j)}]}\right)$. This property is leveraged in the proof of Lem. 2.

To prove the forward direction, we must show that $P^{\mathcal{M}}(\mathbf{F} = \mathbf{f}) = P^{\mathcal{M}'}(\mathbf{F} = \mathbf{f})$ for all $\mathbf{f} \in \mathcal{D}_{\mathbf{F}}$ implies $\mathcal{L}_3(\mathcal{M}) = \mathcal{L}_3(\mathcal{M}')$. Consider any arbitrary \mathcal{L}_3 query from \mathcal{M} ,

$$Q = P^{\mathcal{M}}(\mathbf{y}_{1[x_1]}, \mathbf{y}_{2[x_2]}, \dots) \in \mathcal{L}_3(\mathcal{M}).$$

Denote Q' as the equivalent value from \mathcal{M}' . Denote \mathbf{Y}_* as the set of all of the counterfactual terms of Q . For each term $\mathbf{y}_{i[x_i]}$, consider its ancestral set $An(\mathbf{y}_{i[x_i]})$, and denote $An(\mathbf{Y}_*)$ as the union of all of these sets. Then note that

$$Q = \sum_{\mathbf{a} \in An(\mathbf{Y}_*) \setminus \mathbf{Y}_*} P^{\mathcal{M}}(An(\mathbf{y}_{1[x_1]}), An(\mathbf{y}_{2[x_2]}), \dots),$$

that is, Q is equal to the joint distribution of its ancestral set with all of the terms not in the original query marginalized away. Then, by Fact 2, we have

$$Q = \sum_{\mathbf{a} \in An(\mathbf{Y}_*) \setminus \mathbf{Y}_*} P^{\mathcal{M}} \left(\bigwedge_{W_{\mathbf{t} \in An(\mathbf{Y}_*)}} W_{\mathbf{pa}_W} = w \right), \quad (58)$$

where w and \mathbf{pa}_w are defined as specified in Fact 2. This ancestral set factorization leaves every term with a single variable under the intervention of its parents. If there are duplicate terms (i.e. $W_{\mathbf{pa}_W}$ appears multiple times for the same W and \mathbf{pa}_W), then either they can be reduced to a single term if every value of w is equal (since $p \wedge p = p$ for any proposition p), or, if w is not equal for every term, then $Q = 0$ for both \mathcal{M} and \mathcal{M}' since $W_{\mathbf{pa}_W}$ cannot be equal to two values at once.

Finally, we note that the probability term is simply a marginalized quantity from the functional counterfactual, so we see that

$$Q = \sum_{\mathbf{a} \in An(\mathbf{Y}_*) \setminus \mathbf{Y}_*} \left(\sum_{\mathbf{f}' \in \mathbf{F} \setminus \mathbf{W}_*} P^{\mathcal{M}}(\mathbf{F} = \mathbf{f}') \right),$$

where \mathbf{W}_* refers to the set of $W_{\mathbf{pa}_W}$ terms in Eq. 58, and \mathbf{f} is defined for each value of $V_{\mathbf{pa}_V} \in \mathbf{F}$ as w if $V_{\mathbf{pa}_V} \in \mathbf{W}_*$ or the value from \mathbf{f}' otherwise.

Therefore, since $P^{\mathcal{M}}(\mathbf{F} = \mathbf{f}) = P^{\mathcal{M}'}(\mathbf{F} = \mathbf{f})$, this implies that $Q = Q'$. With this being true for all values of $Q \in \mathcal{L}_3(\mathcal{M})$, this means that $\mathcal{L}_3(\mathcal{M}) = \mathcal{L}_3(\mathcal{M}')$. \square

A.2 Proofs of Sec. 2

The abstractions in this work follow the theory developed by Beckers and Halpern (2019). We first note that Beckers and Halpern (2019) utilizes the idea of ‘‘allowed interventions’’. Specifically, for an SCM \mathcal{M}_L over variables \mathbf{V}_L , the notation \mathcal{I}_L is used to indicate a set of interventions on \mathbf{V}_L that are ‘‘allowed’’ in \mathcal{M}_L . This is relevant when defining interventions across abstractions, since not all interventions on the lower level will have a corresponding intervention on the higher level, as will be shown. Further, the notation \mathcal{I}_L^* is used to define the set of all possible interventions over \mathbf{V}_L .

We use the following definitions, translated to use the notation in our work¹⁵.

¹⁵Note that there can be at most one such possible intervention $\mathbf{X}_H \leftarrow \mathbf{x}_H$ such that $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L) = \mathbf{X}_H \leftarrow \mathbf{x}_H$. It is possible that no such intervention exists, but for this work, we only consider cases where there $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L)$ exists for all $\mathbf{X}_L \leftarrow \mathbf{x}_L \in \mathcal{I}_L$.

Definition 14 ((Beckers and Halpern 2019, Def. 3.12)). Some relevant notation is defined as follows:

- Given a set of variables $\mathbf{V}, \mathbf{X} \subseteq \mathbf{V}$, and $\mathbf{x} \in \mathcal{D}_{\mathbf{X}}$, let $\text{Rst}(\mathbf{V}, \mathbf{x}) = \{\mathbf{v} \in \mathcal{D}_{\mathbf{V}} : \mathbf{v} \text{ is consistent with } \mathbf{x}\}$.
- Given variables \mathbf{V}_L and \mathbf{V}_H , mapping $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$, and value set $\mathbf{T} \subseteq \mathcal{D}_{\mathbf{V}_L}$, denote $\tau(\mathbf{T}) = \{\tau(\mathbf{v}_L) : \mathbf{v}_L \in \mathbf{T}\}$.
- Given allowed interventions \mathcal{I}_L and \mathcal{I}_H over \mathbf{V}_L and \mathbf{V}_H respectively, define $\omega_\tau : \mathcal{I}_L \rightarrow \mathcal{I}_H$ such that $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L) = \mathbf{X}_H \leftarrow \mathbf{x}_H$, where $\tau(\text{Rst}(\mathbf{V}_L, \mathbf{x}_L)) = \text{Rst}(\mathbf{V}_H, \mathbf{x}_H)$.

Definition 15 (τ -Abstraction (Beckers and Halpern 2019, Def. 3.13)). Let $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle$ and $\mathcal{M}_H = \langle \mathbf{U}_H, \mathbf{V}_H, \mathcal{F}_H, P(\mathbf{U}_H) \rangle$ be two SCMs. Let \mathcal{I}_L and \mathcal{I}_H be the sets of allowed interventions respectively. Given $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$, we say that $(\mathcal{M}_H, \mathcal{I}_H)$ is a τ -abstraction of $(\mathcal{M}_L, \mathcal{I}_L)$ if:

1. τ is surjective;
2. There exists surjective $\tau_{\mathbf{U}} : \mathcal{D}_{\mathbf{U}_L} \rightarrow \mathcal{D}_{\mathbf{U}_H}$ that is compatible with τ , i.e.

$$\tau(\mathcal{M}_{L[\mathbf{X}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}_L)) = \mathcal{M}_{H[\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}_L)), \quad (59)$$

for all $\mathbf{u}_L \in \mathcal{D}_{\mathbf{U}_L}$ and all $(\mathbf{X}_L \leftarrow \mathbf{x}_L) \in \mathcal{I}_L$;

3. $\mathcal{I}_H = \omega_\tau(\mathcal{I}_L)$. \blacksquare

Further, we will assume that if $(\mathcal{M}_H, \mathcal{I}_H)$ is a τ -abstraction of $(\mathcal{M}_L, \mathcal{I}_L)$, then $P(\mathbf{U}_H) = \tau_{\mathbf{U}}(P(\mathbf{U}_L)) = P(\tau_{\mathbf{U}}(\mathbf{U}_L))$, that is, the distribution of $P(\mathbf{U}_H)$ can be obtained from $P(\mathbf{U}_L)$ via the push-forward measure through $\tau_{\mathbf{U}}$. While it is not explicitly stated in the definition, this property aligns with the intention of linking the spaces of \mathbf{U}_L and \mathbf{U}_H through $\tau_{\mathbf{U}}$.

Definition 16 (Strong τ -Abstraction (Beckers and Halpern 2019, Def. 3.15)). We say that \mathcal{M}_H is a strong τ -abstraction of \mathcal{M}_L if $(\mathcal{M}_H, \mathcal{I}_H)$ is a τ -abstraction of $(\mathcal{M}_L, \mathcal{I}_L)$ and $\mathcal{I}_H = \mathcal{I}_L^*$. \blacksquare

Definition 17 (Constructive τ -Abstraction (Beckers and Halpern 2019, Def. 3.19)). \mathcal{M}_H is a constructive τ -abstraction of \mathcal{M}_L if \mathcal{M}_H is a strong τ -abstraction of \mathcal{M}_L , and there exists a partition of $\mathbf{V}_L, \mathbb{C} = \{\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_{n+1}\}$ (where $n = |\mathbf{V}_H|$) with nonempty \mathbf{C}_1 to \mathbf{C}_n , such that τ can be decomposed as $\tau = (\tau_{\mathbf{C}_1}, \tau_{\mathbf{C}_2}, \dots, \tau_{\mathbf{C}_n})$, where each $\tau_{\mathbf{C}_i} : \mathcal{D}_{\mathbf{C}_i} \rightarrow \mathcal{D}_{\mathbf{V}_{H,i}}$ maps the i th partition to the i th variable of \mathbf{V}_H . \blacksquare

In typical causal inference settings where nature is modeled by an SCM, every possible intervention is well-defined. In practice, some interventions may not intuitively correspond to an explicit action. For example, in a medical dataset, perhaps cholesterol level is measured from each person in the dataset. It may not make sense to consider interventions on cholesterol level, since it is not clear how, in practice, one would fix or change someone’s cholesterol level to specific values. Nonetheless, it is still possible to theoretically study the effects of such an intervention through the semantics of SCMs. Therefore, the notion of ‘‘allowed interventions’’ is

not typically discussed outside of works of abstractions as every intervention is well-defined.

This no longer holds true in the discussion of abstractions, which is why definitions of abstractions like Def. 15 explicitly require the allowed interventions to be specified. Notably, since every intervention in an SCM involves the creation of a submodel through the mutilation procedure, an intervention on a lower level SCM may not have an obvious counterpart in the higher level SCM. For an in depth understanding of the intricacies of this, we defer readers to read Beckers and Halpern (2019). In this work, given that one would like to use the higher level SCM \mathcal{M}_H for downstream inference tasks, we provide the maximum possible flexibility and assume that $\mathcal{I}_H = \mathcal{I}_H^*$.

Defining \mathcal{I}_L is somewhat trickier, since in any nontrivial abstraction, there are cases where an intervention on the lower level does not have an equivalent intervention on the higher level. To decide on a criteria of which interventions are allowed on the lower level model, consider the following lemmas.

Lemma 3. *Let τ be a constructive abstraction function w.r.t. \mathbb{C} and \mathbb{D} . If \mathbf{X}_L is a union of clusters in \mathbb{C} (that is, there exists $\mathbb{C}' \subseteq \mathbb{C}$ such that $\mathbf{X}_L = \bigcup_{\mathbf{C}_i \in \mathbb{C}'} \mathbf{C}_i$), then $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L)$ exists and is equal to $\tau(\mathbf{X}_L) \leftarrow \tau(\mathbf{x}_L)$. ■*

Proof. If \mathbf{X}_L is a union of clusters in \mathbb{C} , then without loss of generality, suppose it can be decomposed as $\mathbf{X}_L = (\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_k)$, and $\mathbf{V}_L \setminus \mathbf{X}_L = (\mathbf{C}_{k+1}, \mathbf{C}_{k+2}, \dots, \mathbf{C}_n)$. By Def. 6, we can then compute, for any $\mathbf{v}_L \in \text{Rst}(\mathbf{V}_L, \mathbf{x}_L)$,

$$\begin{aligned} \tau(\mathbf{v}_L) &= \tau((\mathbf{x}_L, \mathbf{v}_L \setminus \mathbf{x}_L)) \\ &= \tau((\mathbf{c}_1, \dots, \mathbf{c}_k, \mathbf{c}_{k+1}, \dots, \mathbf{c}_n)) \\ &= (\tau_{\mathbf{C}_1}(\mathbf{c}_1), \dots, \tau_{\mathbf{C}_k}(\mathbf{c}_k), \tau_{\mathbf{C}_{k+1}}(\mathbf{c}_{k+1}), \dots, \tau_{\mathbf{C}_n}(\mathbf{c}_n)) \\ &= (\tau(\mathbf{x}_L), \tau(\mathbf{v}_L \setminus \mathbf{x}_L)) \in \text{Rst}(\tau(\mathbf{V}_L), \tau(\mathbf{x}_L)). \end{aligned}$$

Moreover, for any $\mathbf{v}_H \in \text{Rst}(\tau(\mathbf{V}_L), \tau(\mathbf{x}_L))$, there exists $\mathbf{v}_L \in \text{Rst}(\mathbf{V}_L, \mathbf{x}_L)$ such that $\tau(\mathbf{v}_L) = \mathbf{v}_H$. Specifically, if $\mathbf{v}_H = (\mathcal{D}_{\mathbf{C}_1}^{j_1}, \mathcal{D}_{\mathbf{C}_2}^{j_2}, \dots, \mathcal{D}_{\mathbf{C}_n}^{j_n})$, then any $\mathbf{v}_L = (\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_n)$ satisfies this relationship if $\mathbf{c}_i \in \mathcal{D}_{\mathbf{C}_i}^{j_i}$ for all i . Hence, by definition, $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L)$ exists and is equal to $\tau(\mathbf{X}_L) \leftarrow \tau(\mathbf{x}_L)$. □

Lemma 4. *Let τ be a constructive abstraction function over \mathbb{C} and \mathbb{D} . For any intervention $\mathbf{X}_H \leftarrow \mathbf{x}_H \in \mathcal{I}_H^*$, there exists \mathbf{X}_L such that \mathbf{X}_L is a union of clusters of \mathbb{C} , and $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L) = \mathbf{X}_H \leftarrow \mathbf{x}_H$. ■*

Proof. Without loss of generality, suppose that $\mathbf{X}_H = \{X_1, X_2, \dots, X_k\}$, corresponding to clusters $\{\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_k\}$ respectively. Then, choose $\mathbf{X}_L = \bigcup_{i=1}^k \mathbf{C}_i$, a union of clusters. The proof holds if for each \mathbf{C}_i , there exist values $\mathbf{c}_i \in \mathcal{D}_{\mathbf{C}_i}$ such that $\tau_{\mathbf{C}_i}(\mathbf{c}_i) = x_i$. Note that by definition, x_i corresponds to some $\mathcal{D}_{\mathbf{C}_i}^j$ such that $\tau_{\mathbf{C}_i}(\mathbf{c}_i) = x_i$ for all $\mathbf{c}_i \in \mathcal{D}_{\mathbf{C}_i}^j$. As $\mathcal{D}_{\mathbf{C}_i}^j$ is an element of a partition of $\mathcal{D}_{\mathbf{C}_i}$, it must be nonempty. Hence, the claim holds, and \mathbf{x}_L can be constructed by taking the one such value \mathbf{c}_i for each i . Lemma 3 can then be used to show that $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L) = \tau(\mathbf{X}_L) \leftarrow \tau(\mathbf{x}_L)$, concluding the proof. □

In other words, by Lemma 3, an intervention on a union of intervariable clusters will always have an intuitive well-defined corresponding intervention. Moreover, Lemma 4 shows that all high level interventions are accounted for on the lower level model¹⁶. Hence, we make the assumption that \mathcal{I}_L is defined such that $\mathbf{X}_L \leftarrow \mathbf{x}_L \in \mathcal{I}_L$ if and only if \mathbf{X}_L is a union of clusters of \mathbb{C} . If, for some reason, this choice of \mathcal{I}_L is not desirable for some application, it may indicate that the choice of clusters, \mathbb{C} and \mathbb{D} , should be revised. Under these assumptions, the work in the main text can be presented without the need to explicitly consider allowed interventions, simplifying the discussion.

We now show the main connection between this work and established works by proving Prop. 1.

Proposition 1 (Abstraction Connection). *Let $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ be a constructive abstraction function (Def. 6). \mathcal{M}_H is \mathcal{L}_3 - τ consistent (Def. 7) with \mathcal{M}_L if and only if there exists SCMs \mathcal{M}'_L and \mathcal{M}'_H s.t. $\mathcal{L}_3(\mathcal{M}'_L) = \mathcal{L}_3(\mathcal{M}_L)$, $\mathcal{L}_3(\mathcal{M}'_H) = \mathcal{L}_3(\mathcal{M}_H)$, and \mathcal{M}'_H is a constructive τ -abstraction of \mathcal{M}'_L . ■*

Proof. For this proof, define $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle$ and $\mathcal{M}_H = \langle \mathbf{U}_H, \mathbf{V}_H, \mathcal{F}_H, P(\mathbf{U}_H) \rangle$.

We first show the forward direction: if \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L , then there exists SCM \mathcal{M}'_H such that $\mathcal{L}_3(\mathcal{M}'_H) = \mathcal{L}_3(\mathcal{M}_H)$ and \mathcal{M}'_H is a constructive τ -abstraction of \mathcal{M}_L .

We will start this proof by first attempting to show that \mathcal{M}_H itself is a constructive τ -abstraction of \mathcal{M}_L . We first note that the intervariable clustering \mathbb{C} is defined to be a partition of a subset of \mathbf{V}_L , and τ is constructed by Def. 6, which is specifically defined to be decomposed into subfunctions $\tau_{\mathbf{C}_i}$ mapping cluster \mathbf{C}_i to a corresponding $V_{H,i}$ for each $i \in \{1, \dots, n\}$. If we further define $\mathbf{C}_{n+1} = \mathbf{V}_L \setminus \bigcup_{\mathbf{C}_i \in \mathbb{C}} \mathbf{C}_i$ (variables that are projected out through the abstraction), then we see that $\{\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_{n+1}\}$ forms a partition over \mathbf{V}_L .

What remains is to show that \mathcal{M}_H is a strong τ -abstraction of \mathcal{M}_L . As $\mathcal{I}_H = \mathcal{I}_H^*$ by assumption, this reduces to showing that it is a τ -abstraction. We show that the three requirements of Def. 15 hold:

1. Consider any $\mathbf{v}_H \in \mathcal{D}_{\mathbf{V}_H}$. For all $V_{H,i} \in \mathbf{V}_H$, note that $V_{H,i}$ must correspond to some \mathbf{C}_i , and $v_{H,i}$ must correspond to some $\mathcal{D}_{\mathbf{C}_i}^j$ by definition. Since $\mathcal{D}_{\mathbf{C}_i}^j$ is part of a partition, it must be nonempty, so $\tau_{\mathbf{C}_i}(\mathbf{c}_i) = v_{H,i}$ for any choice of $\mathbf{c}_i \in \mathcal{D}_{\mathbf{C}_i}^j$. Hence, $\tau_{\mathbf{C}_i}$ is surjective for all i , implying that τ is surjective as it is simply a collection of all $\tau_{\mathbf{C}_i}$.
2. Consider the functional counterfactual set \mathbf{F}_H of \mathcal{M}_H (as defined in Eq. 57). For every $\mathbf{f}_H \in \mathcal{D}_{\mathbf{F}_H}$, consider the

¹⁶There are some contrived settings in which $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L)$ still exists even when \mathbf{X}_L is not a union of clusters, but it is inconsequential to omit these cases since all possible high level interventions are covered by Lemma 4.

counterfactual quantity

$$P(\mathbf{F}_H = \mathbf{f}_H) = P\left(\bigwedge_{V_{H,i} \in \mathbf{V}_H} \bigwedge_{\mathbf{pa}_{V_{H,i}}^{(j)} \in \mathcal{D}_{\mathbf{Pa}_{V_{H,i}}}} V_{H,i}[\mathbf{pa}_{V_{H,i}}^{(j)}] = v_{H,i}^{(j)}\right).$$

By Eq. 10, this is equal to

$$P(\mathbf{F}_H = \mathbf{f}_H) = \sum_{\forall i,j \mathbf{c}_{L,i}^{(j)} \in \mathcal{D}_{\mathbf{V}_L} : \tau(\mathbf{c}_{L,i}^{(j)}) = v_{H,i}^{(j)}} P\left(\bigwedge_{i,j} \mathbf{C}_{L,i}[\mathbf{x}_{L,i}^{(j)}] = \mathbf{c}_{L,i}^{(j)}\right)$$

for all choices of $\mathbf{x}_{L,i}^{(j)}$ such that $\tau(\mathbf{x}_{L,i}^{(j)}) = \mathbf{pa}_{V_{H,i}}^{(j)}$. Recall that $\mathcal{D}_{\mathbf{U}_H}(\mathbf{f}_H) \subseteq \mathcal{D}_{\mathbf{U}_H}$ is the set of values of \mathbf{U}_H such that $\mathbf{u}_H \in \mathcal{D}_{\mathbf{U}_H}(\mathbf{f}_H)$ if and only if $\mathbf{F}_H = \mathbf{f}_H$ when $\mathbf{U}_H = \mathbf{u}_H$. Similarly, define $\mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H) \subseteq \mathcal{D}_{\mathbf{U}_L}$ as the set of values of \mathbf{U}_L such that $\mathbf{u}_L \in \mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H)$ if and only if $\bigwedge_{i,j} \mathbf{C}_{L,i}[\mathbf{x}_{L,i}^{(j)}] = \mathbf{c}_{L,i}^{(j)}$ when $\mathbf{U}_L = \mathbf{u}_L$.

Note that there may exist $\mathbf{u}_L \in \mathcal{D}_{\mathbf{U}_L}$ that do not belong to $\mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H)$ for any choice of \mathbf{f}_H . However, the total probability measure of all such cases must be 0, since $P(\mathbf{U}_L \in \mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H)) = P(\mathbf{U}_H \in \mathcal{D}_{\mathbf{U}_H}(\mathbf{f}_H))$, and $\{\mathcal{D}_{\mathbf{U}_H}(\mathbf{f}_H) : \mathbf{f}_H \in \mathcal{D}_{\mathbf{F}_H}, \mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H) \neq \emptyset\}$ forms a partition over $\mathcal{D}_{\mathbf{U}_H}$ by Lemma 1 and must therefore have probability 1. Hence, we can define \mathbf{U}'_L with domain $\mathcal{D}_{\mathbf{U}'_L}$ that is equivalent to $\mathcal{D}_{\mathbf{U}_L}$ but with these cases omitted. Correspondingly, we can define \mathcal{F}' as \mathcal{F} but excluding the outputs when \mathbf{U}_L takes a value not in $\mathcal{D}_{\mathbf{U}'_L}$. We can then define $\mathcal{M}'_L = \langle \mathbf{U}'_L, \mathbf{V}_L, \mathcal{F}'_L, P(\mathbf{U}'_L) \rangle$, where $\mathcal{L}_3(\mathcal{M}'_L) = \mathcal{L}_3(\mathcal{M}_L)$ because only a measure zero portion of $P(\mathbf{U}_L)$ was removed.

Define U' as a random variable with domain $\mathcal{D}_{U'} = \{u'_{\mathbf{f}_H} : \mathbf{f}_H \in \mathcal{D}_{\mathbf{F}_H}, \mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H) = \emptyset\}$, that is, a value for each choice of \mathbf{f}_H with a nonempty $\mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H)$. Define $P(U' = u'_{\mathbf{f}_H}) = P(\mathbf{F}_H = \mathbf{f}_H)$. Choose $\tau_U : \mathcal{D}_{U'} \rightarrow \mathcal{D}_{\mathbf{U}_L}$ such that $\tau_U(u'_L) = u'_L$ if and only if $u'_L \in \mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H)$. Note that this function is surjective because $u'_{\mathbf{f}_H} \in \mathcal{D}_{U'}$ only if $\mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H)$ is nonempty. Furthermore, $P(U' = u'_{\mathbf{f}_H}) = P(\mathbf{F}_H = \mathbf{f}_H) = P(\mathbf{u}_L \in \mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H))$, so the probability distributions are consistent.

It is not necessarily the case that property 2 of τ -abstractions holds between \mathcal{M}_H and \mathcal{M}_L , but we can create a new SCM $\mathcal{M}'_H = \langle U', \mathbf{V}_H, \mathcal{F}'_H, P(U') \rangle$ such that the property holds between \mathcal{M}'_H and \mathcal{M}'_L . Define $\mathcal{F}'_H = \{f'_{V_{H,i}} : V_{H,i} \in \mathbf{V}_H\}$ such that each $f'_{V_{H,i}}(\mathbf{pa}_{V_{H,i}}, u'_{\mathbf{f}_H}) = f_{V_{H,i}}(\mathbf{pa}_{V_{H,i}}, \mathbf{u}_H)$ if and only if $\mathbf{u}_H \in \mathcal{D}_{\mathbf{U}_H}(\mathbf{f}_H)$. Given that $P(U' = u'_{\mathbf{f}_H}) = P(\mathbf{F}_H = \mathbf{f}_H) = P(\mathbf{u}_H \in \mathcal{D}_{\mathbf{U}_H}(\mathbf{f}_H))$, this implies that $P^{\mathcal{M}'_H}(\mathbf{F}_H = \mathbf{f}_H) = P^{\mathcal{M}_H}(\mathbf{F}_H = \mathbf{f}_H)$, further implying that $\mathcal{L}_3(\mathcal{M}'_H) = \mathcal{L}_3(\mathcal{M}_H)$ by Lemma 2.

We now show that $\tau(\mathcal{M}'_{L[\mathbf{x}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}'_L)) = \mathcal{M}'_{H[\omega_\tau(\mathbf{x}_L \leftarrow \mathbf{x}_L)]}(\tau_U(\mathbf{u}'_L))$, for all $\mathbf{u}'_L \in \mathcal{D}_{U'}$ and

all $(\mathbf{X}_L \leftarrow \mathbf{x}_L) \in \mathcal{I}_L$. For the rest of the proof, assume that any notation involving subscripts L and H refers to \mathcal{M}'_L and \mathcal{M}'_H rather than \mathcal{M}_L and \mathcal{M}_H . By Lemma 3, we know that $\omega(\mathbf{X}_L \leftarrow \mathbf{x}_L) = \tau(\mathbf{X}_L) \leftarrow \tau(\mathbf{x}_L)$, so let $\mathbf{x}_H = \tau(\mathbf{x}_L)$. Assume on the contrary that there exists $\mathbf{u}'_L \in \mathcal{D}_{U'}$ and $(\mathbf{X}_L \leftarrow \mathbf{x}_L) \in \mathcal{I}_L$ such that this claim does not hold. Then there must exist $\mathbf{C}_{L,i}$ (and $V_{H,i} = \tau(\mathbf{C}_{L,i})$) such that $\tau(\mathbf{C}_{L,i[\mathbf{x}_L]}(\mathbf{u}'_L)) \neq V_{H,i[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L))$. Let \mathbf{f}_H be the value of \mathbf{F}_H such that this value of $\mathbf{u}'_L \in \mathcal{D}_{\mathbf{U}_L}(\mathbf{f}_H)$. We show using proof by induction that this poses a contradiction.

There must be a topological ordering to \mathbf{V}_H as \mathcal{M}'_H is recursive. In the base case, assume that $\mathbf{Pa}_{V_{H,i}} = \emptyset$. This means that $\tau(\mathbf{C}_{L,i[\mathbf{x}_L]}(\mathbf{u}'_L)) = \tau(\mathbf{C}_{L,i[\emptyset]}(\mathbf{u}'_L))$ and $V_{H,i[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L)) = V_{H,i[\emptyset]}(\tau_U(\mathbf{u}'_L))$. Note that $\tau(\mathbf{C}_{L,i[\emptyset]}(\mathbf{u}'_L)) = \tau(\mathbf{c}_{L,i}^{(j)}) = v_{H,i}^{(j)} = V_{H,i[\emptyset]}(\tau_U(\mathbf{u}'_L))$ by definition of \mathbf{f}_H , contradicting the claim that $\tau(\mathbf{C}_{L,i[\mathbf{x}_L]}(\mathbf{u}'_L)) \neq V_{H,i[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L))$.

Now assume for the sake of induction that $\tau(\mathbf{C}_{L,i'[\mathbf{x}_L]}(\mathbf{u}'_L)) = V_{H,i'[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L))$ for all $V_{H,i'} \in \mathbf{Pa}_{V_{H,i}}$. Note that $\tau(\mathbf{C}_{L,i[\mathbf{x}_L]}(\mathbf{u}'_L)) = \tau(\mathbf{C}_{L,i[\mathbf{Pa}_{\mathbf{C}_{L,i}[\mathbf{x}_L]}(\mathbf{u}'_L)]}(\mathbf{u}'_L))$ and $V_{H,i[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L)) = V_{H,i[\mathbf{Pa}_{V_{H,i}[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L))]}(\tau_U(\mathbf{u}'_L))$ due to Corol. 3.

However, $\tau(\mathbf{Pa}_{\mathbf{C}_{L,i[\mathbf{x}_L]}(\mathbf{u}'_L)}) = \mathbf{Pa}_{V_{H,i[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L))}$ by the inductive hypothesis. This means that $\tau(\mathbf{C}_{L,i[\mathbf{Pa}_{\mathbf{C}_{L,i[\mathbf{x}_L]}(\mathbf{u}'_L)]}(\mathbf{u}'_L)) = \tau(\mathbf{c}_{L,i}^{(j)}) = v_{H,i}^{(j)} = V_{H,i[\mathbf{Pa}_{V_{H,i}[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L))]}(\tau_U(\mathbf{u}'_L))$ from \mathbf{f}_H , once again contradicting the claim that $\tau(\mathbf{C}_{L,i[\mathbf{x}_L]}(\mathbf{u}'_L)) \neq V_{H,i[\mathbf{x}_H]}(\tau_U(\mathbf{u}'_L))$.

Therefore, it must be the case that $\tau(\mathcal{M}'_{L[\mathbf{x}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}'_L)) = \mathcal{M}'_{H[\omega_\tau(\mathbf{x}_L \leftarrow \mathbf{x}_L)]}(\tau_U(\mathbf{u}'_L))$, for all $\mathbf{u}'_L \in \mathcal{D}_{U'}$ and all $(\mathbf{X}_L \leftarrow \mathbf{x}_L) \in \mathcal{I}_L$.

3. $\mathcal{I}_H = \omega_\tau(\mathcal{I}_L)$ holds by the assumption of \mathcal{I}_L and Lemma 4.

This completes the forward direction of the proof.

We now show the backward direction: if there exists SCMs \mathcal{M}'_L and \mathcal{M}'_H such that $\mathcal{L}_3(\mathcal{M}'_L) = \mathcal{L}_3(\mathcal{M}_L)$, $\mathcal{L}_3(\mathcal{M}'_H) = \mathcal{L}_3(\mathcal{M}_H)$, and \mathcal{M}'_H is a constructive τ -abstraction of \mathcal{M}'_L , then \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L . It is sufficient to simply show that \mathcal{M}'_H is \mathcal{L}_3 - τ consistent with \mathcal{M}'_L , since if $\mathcal{L}_3(\mathcal{M}'_L) = \mathcal{L}_3(\mathcal{M}_L)$ and $\mathcal{L}_3(\mathcal{M}'_H) = \mathcal{L}_3(\mathcal{M}_H)$, this would also imply that \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L .

This can be proven by showing that Eq. 10 holds, that is

$$\begin{aligned} & \sum_{\forall i \mathbf{y}_{L,i} \in \mathcal{D}_{\mathbf{Y}_{L,i}} : \tau(\mathbf{y}_{L,i}) = \mathbf{y}_{H,i}} P^{\mathcal{M}'_L}(\mathbf{y}_{L,1[\mathbf{x}_{L,1}]}, \mathbf{y}_{L,2[\mathbf{x}_{L,2}]}, \dots) \\ & = P^{\mathcal{M}'_H}(\mathbf{y}_{H,1[\tau(\mathbf{x}_{H,1})]}, \mathbf{y}_{H,2[\tau(\mathbf{x}_{L,2})]}, \dots) \end{aligned}$$

for all choices of $\mathbf{y}_{H,i}$ and $\mathbf{x}_{H,i}$. Denote Q as the l.h.s. of the equation and $\tau(Q)$ as the r.h.s. Denote $\mathcal{D}_Q \subseteq \mathcal{D}_{U'_L}$ as the set of values of \mathbf{u}'_L such that $\bigwedge_i \mathbf{Y}_{L,i[\mathbf{x}_{L,i}]} = \mathbf{y}_{L,i[\mathbf{x}_{L,i}]}$ for all $\mathbf{y}_{L,i}$ such that $\tau(\mathbf{y}_{L,i}) = \mathbf{y}_{H,i}$. Similarly, denote $\mathcal{D}_{\tau(Q)} \subseteq \mathcal{D}_{U'_H}$ as the set of values of \mathbf{u}'_H such that $\bigwedge_i \mathbf{Y}_{H,i[\tau(\mathbf{x}_{L,i})]} = \mathbf{y}_{H,i[\tau(\mathbf{x}_{L,i})]}$. Note that $Q = P(\mathbf{U}'_L \in \mathcal{D}_Q)$ and $\tau(Q) =$

$P(\mathbf{U}'_H \in \mathcal{D}_{\tau(Q)})$. We claim that $\mathbf{u}'_L \in \mathcal{D}_Q$ if and only if $\tau_{\mathbf{U}}(\mathbf{u}'_L) \in \mathcal{D}_{\tau(Q)}$.

By definition of constructive τ -abstractions, there must exist $\tau_{\mathbf{U}}$ such that $\tau(\mathcal{M}_{L[\mathbf{x}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}'_L)) = \mathcal{M}_{H[\omega_{\tau}(\mathbf{x}_L \leftarrow \mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}'_L))$, for all $\mathbf{u}'_L \in \mathcal{D}_{\mathbf{U}'_L}$ and all $(\mathbf{X}_L \leftarrow \mathbf{x}_L) \in \mathcal{I}_L$, and further that $P(\mathbf{U}'_H) = \dot{P}(\tau(\mathbf{U}'_L))$ by assumption. This implies that $\tau(\mathbf{Y}_{L,i[\mathbf{x}_L,i]}(\mathbf{u}'_L)) = \mathbf{Y}_{H,i[\omega_{\tau}(\mathbf{x}_L,i)]}(\tau_{\mathbf{U}}(\mathbf{u}'_L)) = \mathbf{Y}_{H,i[\tau(\mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}'_L))$ for all i by Lemma 3. Hence, if $\mathbf{Y}_{L,i[\mathbf{x}_L,i]}(\mathbf{u}'_L) = \mathbf{y}_{L,i}$, where $\tau(\mathbf{y}_{L,i}) = \mathbf{y}_{H,i}$, then $\mathbf{Y}_{H,i[\tau(\mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}'_L)) = \mathbf{y}_{H,i}$. Considering this for all values of i , it must be the case that $\mathbf{u}'_L \in \mathcal{D}_Q$ if and only if $\tau_{\mathbf{U}}(\mathbf{u}'_L) \in \mathcal{D}_{\tau(Q)}$.

Since $P(\mathbf{U}'_H) = P(\tau(\mathbf{U}'_L))$, this implies that $\tau(Q) = P(\mathbf{U}'_H \in \mathcal{D}_{\tau(Q)}) = P(\mathbf{U}'_L \in \mathcal{D}_Q) = Q$, concluding the proof. \square

We note that the theorem does not claim \mathcal{L}_3 - τ consistency is equivalent to constructive τ -abstractions, rather making a weaker claim that there must exist a pair of (potentially different) SCMs that are \mathcal{L}_3 -equivalent and fit the definition of a constructive τ -abstraction. The reason is that, in fact, the definition for constructive τ -abstraction is stricter than \mathcal{L}_3 - τ consistency, but as evident in the proof, the only restriction is on the domains of \mathbf{U}_L and \mathbf{U}_H . In the case of \mathbf{U}_L , there may be a measure zero portion of the domain $\mathcal{D}_{\mathbf{U}_L}$ that do not translate to the higher level functional counterfactuals, and in the case of \mathbf{U}_H , it is possible that the space of $\mathcal{D}_{\mathbf{U}_H}$ may not allow for $\tau_{\mathbf{U}}$ to be surjective. However, the theorem still essentially states that they are equivalent, at least on the three levels of the PCH. We therefore argue that, given the unobserved nature of the exogenous variables and the generating SCM, these two concepts are equivalent on a practical level.

We now prove that Alg. 1 successfully returns an \mathcal{L}_3 - τ consistent model.

Proposition 3. *Let τ and \mathcal{M}_H be the function and SCM obtained from running Alg. 1 on inputs \mathcal{M}_L , \mathbb{C} , and \mathbb{D} . Then, \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L .* \blacksquare

Proof. We can show this result by first showing that the output of Alg. 1, \mathcal{M}_H , is a constructive τ -abstraction of \mathcal{M}_L . By Def. 6, it is clear by construction that τ can be decomposed as $\tau = (\tau_{\mathbf{C}_1}, \tau_{\mathbf{C}_2}, \dots, \tau_{\mathbf{C}_n})$, and each $\tau_{\mathbf{C}_i}$ maps the i th partition to the i th variable of \mathbf{V}_H , as established on line 2. Hence, we must simply show that \mathcal{M}_H is a τ -abstraction of \mathcal{M}_L . Out of the three properties of Def. 15, properties 1 and 3 can be proven similarly to how it is done in the forward direction of Prop. 1. Then, to prove property 2, we must simply show that there exists surjective $\tau_{\mathbf{U}} : \mathcal{D}_{\mathbf{U}_L} \rightarrow \mathcal{D}_{\mathbf{U}_H}$ such that $P(\mathbf{U}_H) = \tau_{\mathbf{U}}(P(\mathbf{U}_L)) = P(\tau_{\mathbf{U}}(\mathbf{U}_L))$, and $\tau(\mathcal{M}_{L[\mathbf{x}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}_L)) = \mathcal{M}_{H[\omega_{\tau}(\mathbf{x}_L \leftarrow \mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}_L))$, for all $\mathbf{u}_L \in \mathcal{D}_{\mathbf{U}_L}$ and all $(\mathbf{X}_L \leftarrow \mathbf{x}_L) \in \mathcal{I}_L$.

The choice of $\tau_{\mathbf{U}}$ is simple—we can use the identity function as by construction in line 1, $\mathbf{U}_H = \mathbf{U}_L$ and $P(\mathbf{U}_H) = P(\mathbf{U}_L)$. Further, note that $\omega_{\tau}(\mathbf{X}_L \leftarrow \mathbf{x}_L) = \tau(\mathbf{X}_L) \leftarrow \tau(\mathbf{x}_L)$ by Lemma 3. Hence, we must simply show that $\tau(\mathcal{M}_{L[\mathbf{x}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}_L)) = \mathcal{M}_{H[\tau(\mathbf{x}_L) \leftarrow \tau(\mathbf{x}_L)]}(\mathbf{u}_L)$.

We start by showing that $V_{i[\tau(\mathbf{x}_L)]}(\mathbf{u}_L) = \tau(\mathbf{C}_{i[\mathbf{x}_L]}(\mathbf{u}_L))$ for every $V_i \in \mathbf{V}_H$ and corresponding $\mathbf{C}_i \in \mathbb{C}$. Note that for every $V_i \in \mathbf{V}_H$, line 5 dictates that

$$V_i \leftarrow f_i^H(\mathbf{pa}_{V_i}, \mathbf{u}_{V_i}) = \tau(f_V^L(\widetilde{\mathbf{pa}}_V, \mathbf{u}_V) : V \in \mathbf{C}_i),$$

which implies that

$$V_{i[\tau(\mathbf{x}_L)]}(\mathbf{u}_L) = \tau(f_V^L(\widetilde{\mathbf{pa}}_V, \mathbf{u}_V) : V \in \mathbf{C}_i), \quad (60)$$

where \mathbf{u}_V is compatible with \mathbf{u}_L and $\widetilde{\mathbf{pa}}_V$ refers to any value $\mathbf{pa}_V \in \mathcal{D}_{\mathbf{Pa}_V}$ such that $\tau(\mathbf{pa}_V) = \mathbf{Pa}_{V_i[\tau(\mathbf{x}_L)]}(\mathbf{u}_L)$. Note that since \mathcal{M}_L satisfies the abstract invariance condition (AIC) w.r.t. τ , the value of f_V^L will not change based on this choice of \mathbf{pa}_V .

We continue using proof by induction. Since \mathcal{M}_H is recursive, this implies there is a topological ordering of the functions of \mathcal{M}_H . In the base case, if $\mathbf{Pa}_{V_i} = \emptyset$, then

$$V_{i[\tau(\mathbf{x}_L)]}(\mathbf{u}_L) = V_{i[\emptyset]}(\mathbf{u}_L) = \tau(f_V^L(\mathbf{u}_V) : V \in \mathbf{C}_i) = \tau(\mathbf{C}_i(\mathbf{u}_L)),$$

which aligns with the claim. For the inductive hypothesis, assume that $V_{j[\tau(\mathbf{x}_L)]}(\mathbf{u}_L) = \tau(\mathbf{C}_{j[\mathbf{x}_L]}(\mathbf{u}_L))$ for every $V_j \in \mathbf{Pa}_{V_i}$. Then we have

$$\begin{aligned} V_{i[\tau(\mathbf{x}_L)]}(\mathbf{u}_L) &= V_{i[\mathbf{Pa}_{V_i[\tau(\mathbf{x}_L)]}(\mathbf{u}_L)]}(\mathbf{u}_L) \\ &\text{by Corol. 3} \\ &= V_{i[\tau(\mathbf{Pa}_{\mathbf{C}_i[\mathbf{x}_L]}(\mathbf{u}_L))]}(\mathbf{u}_L) \\ &\text{by inductive hypothesis} \\ &= \tau(f_V^L(\mathbf{Pa}_{\mathbf{C}_i[\mathbf{x}_L]}(\mathbf{u}_L), \mathbf{u}_V) : V \in \mathbf{C}_i) \\ &\text{by Eq. 60} \\ &= \tau(\mathbf{C}_{i[\mathbf{x}_L]}(\mathbf{u}_L)) \\ &\text{simplification,} \end{aligned}$$

proving the claim. Since $V_{i[\tau(\mathbf{x}_L)]}(\mathbf{u}_L) = \tau(\mathbf{C}_{i[\mathbf{x}_L]}(\mathbf{u}_L))$ for every $V_i \in \mathbf{V}_H$, this implies that $\tau(\mathcal{M}_{L[\mathbf{x}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}_L)) = \mathcal{M}_{H[\omega_{\tau}(\mathbf{x}_L \leftarrow \mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}_L))$, proving that \mathcal{M}_H is a constructive τ -abstraction of \mathcal{M}_L . Finally, Prop. 1 proves that \mathcal{M}_H must therefore be an \mathcal{L}_3 - τ abstraction of \mathcal{M}_L . \square

We leverage this property to prove that the abstract invariance condition (AIC) is necessary and sufficient for the existence of an abstraction.

Proposition 2 (Abstraction Conditions). *For any SCM \mathcal{M}_L and constructive abstraction function τ relative to \mathbb{C} and \mathbb{D} , there exists an SCM \mathcal{M}_H over variables $\mathbf{V}_H = \tau(\mathbf{V}_L)$ such that \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L if and only if there exists \mathcal{M}'_L such that $\mathcal{L}_3(\mathcal{M}_L) = \mathcal{L}_3(\mathcal{M}'_L)$ and \mathcal{M}'_L satisfies the abstract invariance condition with respect to τ .* \blacksquare

Proof. If \mathcal{M}'_L satisfies the AIC w.r.t. τ , then one can use Alg. 1 to obtain an example of a model that is \mathcal{L}_3 - τ consistent with \mathcal{M}'_L (and therefore with \mathcal{M}_L), as proven in Prop. 3.

We now consider the other direction. If \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L , then by Prop. 1, there must exist some \mathcal{M}'_L and \mathcal{M}'_H such that $\mathcal{L}_3(\mathcal{M}'_L) = \mathcal{L}_3(\mathcal{M}_L)$, $\mathcal{L}_3(\mathcal{M}'_H) = \mathcal{L}_3(\mathcal{M}_H)$ and \mathcal{M}'_H is a constructive τ -abstraction of \mathcal{M}'_L . For the rest of this proof, assume that any terms with a subscript of L or H refer to \mathcal{M}'_L

or \mathcal{M}'_H instead of \mathcal{M}_L and \mathcal{M}_H . The constructive τ -abstraction property implies that there exists surjective $\tau_{\mathbf{U}} : \mathcal{D}_{\mathbf{U}'_L} \rightarrow \mathcal{D}_{\mathbf{U}_H}$ such that $\tau(\mathcal{M}'_{L[\mathbf{X}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}'_L)) = \mathcal{M}'_{H[\omega_{\tau}(\mathbf{x}_L \leftarrow \mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}'_L))$, for all $\mathbf{u}'_L \in \mathcal{D}_{\mathbf{U}'_L}$ and all $(\mathbf{X}_L \leftarrow \mathbf{x}_L) \in \mathcal{I}_L$. By Lemma 3, this means $\tau(\mathcal{M}'_{L[\mathbf{X}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}'_L)) = \mathcal{M}'_{H[\tau(\mathbf{X}_L) \leftarrow \tau(\mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}'_L))$.

Assume for the sake of contradiction that it is not the case that \mathcal{M}'_L satisfies the AIC w.r.t. τ . Then there must exist $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{D}_{\mathbf{V}_L}$ such that $\tau(\mathbf{v}_1) = \tau(\mathbf{v}_2)$, yet

$$\begin{aligned} & \tau \left(\left(f_{V_L}^L(\mathbf{pa}_{V_L}^{(1)}, \mathbf{u}'_{V_L}) : V \in \mathbf{C}_i \right) \right) \\ & \neq \tau \left(\left(f_{V_L}^L(\mathbf{pa}_{V_L}^{(2)}, \mathbf{u}'_{V_L}) : V \in \mathbf{C}_i \right) \right) \end{aligned}$$

for some value of $\mathbf{u}'_L \in \mathcal{D}_{\mathbf{U}'_L}$ and $\mathbf{C}_i \in \mathbb{C}$.

This would imply that $\tau \left(\mathbf{C}_{i[\mathbf{pa}_{\mathbf{C}_i}^{(1)}}(\mathbf{u}'_L) \right) \neq \tau \left(\mathbf{C}_{i[\mathbf{pa}_{\mathbf{C}_i}^{(2)}}(\mathbf{u}'_L) \right)$, so $\tau \left(\mathcal{M}'_{L[\mathbf{pa}_{\mathbf{C}_i}^{(1)}}(\mathbf{u}'_L) \right) \neq \tau \left(\mathcal{M}'_{L[\mathbf{pa}_{\mathbf{C}_i}^{(2)}}(\mathbf{u}'_L) \right)$. By the τ -abstraction definition, we have $\tau \left(\mathcal{M}'_{L[\mathbf{pa}_{\mathbf{C}_i}^{(1)}}(\mathbf{u}'_L) \right) = \mathcal{M}'_{H[\tau(\mathbf{pa}_{\mathbf{C}_i}^{(1)})]}(\tau_{\mathbf{U}}(\mathbf{u}'_L))$ and $\tau \left(\mathcal{M}'_{L[\mathbf{pa}_{\mathbf{C}_i}^{(2)}}(\mathbf{u}'_L) \right) = \mathcal{M}'_{H[\tau(\mathbf{pa}_{\mathbf{C}_i}^{(2)})]}(\tau_{\mathbf{U}}(\mathbf{u}'_L))$. However, since $\tau(\mathbf{pa}_{\mathbf{C}_i}^{(1)}) = \tau(\mathbf{pa}_{\mathbf{C}_i}^{(2)})$, this implies that $\tau \left(\mathcal{M}'_{L[\mathbf{pa}_{\mathbf{C}_i}^{(1)}}(\mathbf{u}'_L) \right) = \tau \left(\mathcal{M}'_{L[\mathbf{pa}_{\mathbf{C}_i}^{(2)}}(\mathbf{u}'_L) \right)$, contradicting the earlier statement. Therefore, \mathcal{M}'_L must satisfy the AIC w.r.t. τ , completing the proof. \square

A.3 Proofs of Sec. 3

We start by noting the impossibility of performing causal inferences without additional assumptions, as implied by the Causal Hierarchy Theorem.

Fact 3 (Causal Hierarchy Theorem (CHT) (Bareinboim et al. 2022, Thm. 1)). *Let Ω^* be the set of all SCMs. We say that Layer j of the causal hierarchy for SCMs collapses to Layer i ($i < j$) relative to $\mathcal{M}^* \in \Omega^*$ if $L_i(\mathcal{M}^*) = L_i(\mathcal{M})$ implies that $L_j(\mathcal{M}^*) = L_j(\mathcal{M})$ for all $\mathcal{M} \in \Omega^*$. Then, with respect to the Lebesgue measure over (a suitable encoding of L_3 -equivalence classes of) SCMs, the subset in which Layer j of SCMs collapses to Layer i is measure zero. \blacksquare*

Given this result, we note that the same principle applies to performing causal inferences across abstractions.

Lemma 5. *Let Ω_L and Ω_H be the space of SCMs defined over \mathbf{V}_L and \mathbf{V}_H respectively, and let $\tau : \mathcal{D}_{\mathbf{V}_H} \rightarrow \mathcal{D}_{\mathbf{V}_L}$ be a constructive abstraction function defined over clusters \mathbb{C} and \mathbb{D} . Let Ω'_L be the subset of Ω_L that satisfies the AIC. Define $\psi_{\tau} : \Omega'_L \rightarrow \Omega_H$ such that $\psi_{\tau}(\mathcal{M}_L) = \mathcal{M}_H$, where $\mathcal{M}_H \in \Omega_H$ is $\mathcal{L}_{3-\tau}$ consistent with \mathcal{M}_L (while there could be many such SCMs, they are all \mathcal{L}_3 -equivalent, so we can arbitrarily choose the output of Alg. 1 on inputs \mathcal{M}_L and τ , which must exist due to Prop. 3). Then, ψ_{τ} is surjective (i.e. $\{\psi(\mathcal{M}_L) : \mathcal{M}_L \in \Omega'_L\} = \Omega_H$). \blacksquare*

Proof. For any SCM $\mathcal{M}_H = \langle \mathbf{U}_H, \mathbf{V}_H, \mathcal{F}_H, P(\mathbf{U}_H) \rangle \in \Omega_H$, one can construct SCM $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle \in \Omega'_L$ such that \mathcal{M}_H is $\mathcal{L}_{3-\tau}$ consistent as \mathcal{M}_L as follows:

1. Choose $\mathbf{U}_L = \mathbf{U}_H$ and $P(\mathbf{U}_L) = P(\mathbf{U}_H)$.
2. For each $V_L \in \mathbf{V}_L$, let $\mathbf{C} \in \mathbb{C}$ be the intervariable cluster such that $V_L \in \mathbf{C}$, and let $V_H = \tau(\mathbf{C})$. Define $\mathbf{Pa}_{V_L} \subseteq \mathbf{V}_L$ as the set of variables such that $\tau(\mathbf{Pa}_{V_L}) = \mathbf{Pa}_{V_H}$. Define $\mathbf{U}_{V_L} \subseteq \mathbf{U}_L$ as \mathbf{U}_{V_H} .
3. For all $\mathbf{C} \in \mathbb{C}$, note that for any pair $V_1, V_2 \in \mathbf{C}$, $\mathbf{Pa}_{V_1} = \mathbf{Pa}_{V_2}$ and $\mathbf{U}_{V_1} = \mathbf{U}_{V_2}$. Denote $V_H = \tau(\mathbf{C})$. For each $V \in \mathbf{C}$, choose $f_{V_L}^L \in \mathcal{F}_L$ arbitrarily such that $\tau(\{f_{V_L}^L(\mathbf{pa}_V, \mathbf{u}_V) : V_L \in \mathbf{C}\}) = f_{V_H}^H(\tau(\mathbf{pa}_{V_L}), \mathbf{u}_{V_L})$. There must exist at least one such setting since for any possible input $\mathbf{pa}_{V_H}, \mathbf{u}_{V_H}$ to $f_{V_H}^H$, there is at least one set of inputs $\mathbf{pa}_{V_L}, \mathbf{u}_{V_L}$ to $f_{V_L}^L$ such that $\tau(\mathbf{pa}_{V_L}) = \mathbf{pa}_{V_H}$ (due to surjectivity of τ) and $\mathbf{u}_{V_L} = \mathbf{u}_{V_H}$.

One can easily verify that running Alg. 1 on this choice of \mathcal{M}_L will return \mathcal{M}_H , implying that \mathcal{M}_H is $\mathcal{L}_{3-\tau}$ consistent with \mathcal{M}_L and that ψ_{τ} is surjective. \square

Proposition 4 (Abstract Causal Hierarchy Theorem (Formal Version)). *Let Ω_L and Ω_H be the space of models defined over \mathbf{V}_L and \mathbf{V}_H respectively, and let Ω'_L be the subset of Ω_L such that the abstract invariance condition holds. Let $\tau : \mathcal{D}_{\mathbf{V}_H} \rightarrow \mathcal{D}_{\mathbf{V}_L}$ be a constructive abstraction function. We say that Layer j of the causal hierarchy for Ω_H τ -collapses to Layer i ($i < j$) relative to $\mathcal{M}_L \in \Omega_L$ if \mathcal{L}_i - τ consistency implies \mathcal{L}_j - τ consistency of \mathcal{M}_H with \mathcal{M}_L for all $\mathcal{M}_H \in \Omega_H$. Then, w.r.t. Lebesgue measure over (a suitable encoding of \mathcal{L}_3 -equivalence classes of) Ω'_L , the subset in which Layer j of Ω_H τ -collapses to Layer i has measure zero. \blacksquare*

Proof. We first show that Layer j of Ω_H τ -collapses to Layer i relative to $\mathcal{M}_L \in \Omega'_L$ if and only if Layer j of Ω_H collapses to Layer i relative to $\psi_{\tau}(\mathcal{M}_L) \in \Omega_H$ (as defined in Lem. 5). If Layer j of Ω_H τ -collapses to Layer i relative to \mathcal{M}_L , then that implies that all SCMs in Ω_H that are \mathcal{L}_i - τ consistent with \mathcal{M}_L are also \mathcal{L}_j - τ consistent, including $\psi_{\tau}(\mathcal{M}_L)$. This is only possible if they are all \mathcal{L}_i - and \mathcal{L}_j -consistent with each other, implying regular collapse relative to $\psi_{\tau}(\mathcal{M}_L)$. Conversely, if Layer j of Ω_H collapses to Layer i relative to $\psi_{\tau}(\mathcal{M}_L)$, then all SCMs in Ω_H that are \mathcal{L}_i -consistent with $\psi_{\tau}(\mathcal{M}_L)$ must also be \mathcal{L}_j -consistent. By definition, $\psi_{\tau}(\mathcal{M}_L)$ is $\mathcal{L}_{3-\tau}$ consistent with \mathcal{M}_L , so this implies that all SCMs in Ω_H that are \mathcal{L}_i - τ consistent with \mathcal{M}_L are also \mathcal{L}_j - τ consistent.

Fact 3 states that the subset of Ω_H in which Layer j of Ω_H collapses to Layer i is measure 0. Hence, the subset of Ω'_L (under the same encoding w.r.t. the set $\{\psi(\mathcal{M}_L) : \mathcal{M}_L \in \Omega'_L\} = \Omega_H$ (as proven in Lem. 5)) in which Layer j of Ω_H τ -collapses to Layer i is also measure 0. \square

As a consequence, causal assumptions are necessary to make causal inferences. For this work, we leverage cluster causal diagrams (C-DAGs), from Def. 9.

For the following proofs, consider the classical definition of identifiability.

Definition 18. Let Ω^* be the space containing all SCMs defined over endogenous variables \mathbf{V} . We say that a causal query Q is identifiable (ID) from the available data \mathbb{Z} and the causal diagram \mathcal{G} if $Q(\mathcal{M}_1) = Q(\mathcal{M}_2)$ for every pair of models $\mathcal{M}_1, \mathcal{M}_2 \in \Omega^*$ such that \mathcal{M}_1 and \mathcal{M}_2 both induce \mathcal{G} and $\mathbb{Z}(\mathcal{M}_1) = \mathbb{Z}(\mathcal{M}_2)$. ■

We can now prove that abstract identification is equivalent to classical identification on the higher level.

Theorem 1 (Dual Abstract ID). Consider a counterfactual query Q over \mathbf{V}_L , a constructive abstraction function τ w.r.t. clusters \mathbb{C} and \mathbb{D} , a C-DAG \mathcal{G}_C , and data \mathbb{Z} from \mathbf{V}_L . Q is τ -ID from \mathcal{G}_C and \mathbb{Z} if and only if $\tau(Q)$ is ID from \mathcal{G}_C and $\tau(\mathbb{Z})$. ■

Proof. Let Ω_L and Ω_H be the space of SCMs defined over \mathbf{V}_L and \mathbf{V}_H respectively, and let $\Omega_L(\mathcal{G}_C)$ and $\Omega_H(\mathcal{G}_C)$ be their corresponding subsets that induce graph \mathcal{G}_C . If Q is τ -ID from \mathcal{G}_C and \mathbb{Z} , then every pair of $\mathcal{M}_L \in \Omega_L(\mathcal{G}_C), \mathcal{M}_H \in \Omega_H(\mathcal{G}_C)$ such that \mathcal{M}_H is \mathbb{Z} - τ consistent with \mathcal{M}_L must have \mathcal{M}_H be Q - τ consistent with \mathcal{M}_L . For all such $\mathcal{M}_H, \mathbb{Z}$ - τ consistency and Q - τ consistency with \mathcal{M}_L implies that \mathcal{M}_H is $\tau(\mathbb{Z})$ -consistent and $\tau(Q)$ -consistent by Def. 7. For any pair $\mathcal{M}_1, \mathcal{M}_2 \in \Omega_H$ that induce \mathcal{G}_C , $\tau(\mathbb{Z})(\mathcal{M}_1) = \tau(\mathbb{Z})(\mathcal{M}_2)$ therefore implies that both \mathcal{M}_1 and \mathcal{M}_2 must be \mathbb{Z} - τ consistent with \mathcal{M}_L and must therefore both be Q - τ consistent, so $\tau(Q)(\mathcal{M}_1) = \tau(Q)(\mathcal{M}_2)$. Hence, $\tau(Q)$ is ID from \mathcal{G}_C and $\tau(\mathbb{Z})$ by Def. 18.

Conversely, if $\tau(Q)$ is ID from \mathcal{G}_C and $\tau(\mathbb{Z})$, then for any $\mathcal{M}_1, \mathcal{M}_2 \in \Omega_H$ that induces \mathcal{G}_C such that $\tau(\mathbb{Z})(\mathcal{M}_1) = \tau(\mathbb{Z})(\mathcal{M}_2)$, it must be the case that $\tau(Q)(\mathcal{M}_1) = \tau(Q)(\mathcal{M}_2)$. For every $\mathcal{M}_L \in \Omega_L(\mathcal{G}_C)$, Prop. 3 states that there exists some $\mathcal{M}_H \in \Omega_H(\mathcal{G}_C)$ that is \mathbb{Z} - τ consistent with \mathcal{M}_L , implying that \mathcal{M}_H is both \mathbb{Z} - τ consistent and Q - τ consistent with \mathcal{M}_L . Since all $\mathcal{M}_H \in \Omega_H(\mathcal{G}_C)$ that match in $\tau(\mathbb{Z})$ must also match in $\tau(Q)$, it must be the case that all such \mathcal{M}_H that are \mathbb{Z} - τ consistent with \mathcal{M}_L must also be Q - τ consistent with \mathcal{M}_L . Hence, by definition, Q is τ -ID from \mathcal{G}_C and \mathbb{Z} . □

We also connect this result to the results of neural identification with NCMs.

Definition 19 (Neural Counterfactual Identification (Xia, Pan, and Bareinboim 2023, Def. 4)). Consider an SCM \mathcal{M}^* and the corresponding causal diagram \mathcal{G} . Let $\mathbb{Z} = \{P(\mathbf{V}_{\mathbf{z}_k})\}_{k=1}^{\ell}$ be a collection of available interventional (or observational if $\mathbf{z}_k = \emptyset$) distributions from \mathcal{M}^* . The counterfactual query $P(\mathbf{Y}_* = \mathbf{y}_* \mid \mathbf{X}_* = \mathbf{x}_*)$ is said to be neural identifiable (identifiable, for short) from the set of \mathcal{G} -constrained NCMs $\Omega(\mathcal{G})$ and \mathbb{Z} if and only if $P^{\widehat{\mathcal{M}}_1}(\mathbf{y}_* \mid \mathbf{x}_*) = P^{\widehat{\mathcal{M}}_2}(\mathbf{y}_* \mid \mathbf{x}_*)$ for every pair of models $\widehat{\mathcal{M}}_1, \widehat{\mathcal{M}}_2 \in \Omega(\mathcal{G})$ s.t. they match \mathcal{M}^* on all distributions in \mathbb{Z} (i.e. $\mathbb{Z}(\mathcal{M}^*) = \mathbb{Z}(\mathcal{M}_1) = \mathbb{Z}(\mathcal{M}_2) > 0$). ■

Fact 4 (Counterfactual Graphical-Neural Equivalence (Dual ID) (Xia, Pan, and Bareinboim 2023, Thm. 3)). Let Ω^*, Ω be the spaces including all SCMs and NCMs, respectively. Consider the true SCM \mathcal{M}^* and the corresponding causal diagram \mathcal{G} . Let $Q = P(\mathbf{y}_* \mid \mathbf{x}_*)$ be the target query and \mathbb{Z} the set of observational and interventional distributions

available. Then, Q is neural identifiable from $\Omega(\mathcal{G})$ and \mathbb{Z} if and only if it is identifiable from \mathcal{G} and \mathbb{Z} . ■

Fact 5 (Neural Counterfactual Mutilation (Operational ID) (Xia, Pan, and Bareinboim 2023, Corol. 1)). Consider the true SCM $\mathcal{M}^* \in \Omega^*$, causal diagram \mathcal{G} , a set of available distributions \mathbb{Z} , and a target query Q equal to $P^{\mathcal{M}^*}(\mathbf{y}_* \mid \mathbf{x}_*)$. Let $\widehat{\mathcal{M}} \in \Omega(\mathcal{G})$ be a \mathcal{G} -constrained NCM such that $\mathbb{Z}(\widehat{\mathcal{M}}) = \mathbb{Z}(\mathcal{M}^*)$. If Q is identifiable from \mathcal{G} and \mathbb{Z} , then Q is computable via Eq. 1 from $\widehat{\mathcal{M}}$. ■

The connection between abstract identification and neural identification follows naturally.

Corollary 1 (Abstract ID with NCMs). Q is τ -ID from \mathcal{G}_C and \mathbb{Z} iff $\tau(Q)$ is Neural-ID from $\widehat{\Omega}(\mathcal{G}_C)$ and $\tau(\mathbb{Z})$. Moreover, if it is ID, then Q can be computed by $\tau(Q)$ from any \mathcal{G}_C -NCM $\widehat{\mathcal{M}}$ that is $\tau(\mathbb{Z})$ -consistent. ■

Proof. This is a direct consequence of Thm. 1, Fact 4, and Fact 5. □

Corollary 2 (Soundness and Completeness). Let \mathcal{M}_L be the low-level SCM, \mathbb{C} and \mathbb{D} be inter/intravariabile clusters of \mathbf{V}_L , \mathcal{G}_C be a C-DAG, Q be a query, and \widehat{Q} be the result from running Alg. 2 with inputs $\mathbb{Z}(\mathcal{M}_L) > 0, \mathbb{C}, \mathbb{D}, \mathcal{G}_C$, and Q . Then, Q is τ -ID from \mathcal{G}_C and \mathbb{Z} if and only if \widehat{Q} is not FAIL. Moreover, if \widehat{Q} is not FAIL, then $\widehat{Q} = Q(\mathcal{M}_L)$. ■

Proof. Lines 1-2 of Alg. 2 constructs τ given \mathbb{C} and \mathbb{D} . Lines 3-9 checks that $\tau(Q)$ is neural identifiable from $\widehat{\Omega}(\mathcal{G}_C)$ and $\tau(\mathbb{Z})$. Lines 4 and 5 find the two parameterizations θ_{\min}^* and θ_{\max}^* that minimize and maximize $\tau(Q)$ while simultaneously guaranteeing $\tau(\mathbb{Z})$ -consistency. Hence, if the two parameterizations result in the same value for $\tau(Q)$, then all such NCMs must match in $\tau(Q)$, guaranteeing neural identifiability. Otherwise, the two parameterizations provide the counterexample for two NCMs that do not match in $\tau(Q)$. Finally, Corol. 1 states that neural identifiability implies abstract identifiability. □

A.4 Proofs of Sec. 4

We first start by showing the following result.

Lemma 6. For any choice of intravariabile clusters \mathbb{D} such that \mathcal{M}_L satisfies the AIC w.r.t. the corresponding τ , \mathcal{M}_L will also satisfy the AIC w.r.t. any finer clustering \mathbb{D}' (i.e. for all $\mathbb{D}_{\mathbf{C}_i} \in \mathbb{D}$ and all $\mathcal{D}_{\mathbf{C}_i}^{(j)} \in \mathbb{D}_{\mathbf{C}_i}, \mathcal{D}_{\mathbf{C}_i}^{(j)}$ is a subset of some $\mathcal{D}_{\mathbf{C}_i}^{(j')} \in \mathbb{D}'_{\mathbf{C}_i}$). ■

Proof. Fix intervariable clusters \mathbb{C} and denote τ and τ' as the constructive abstraction function defined w.r.t. (\mathbb{C}, \mathbb{D}) and $(\mathbb{C}, \mathbb{D}')$ respectively. If \mathcal{M}_L satisfies that AIC w.r.t. τ , that implies that for all $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{D}_{\mathbf{V}_L}$ such that $\tau(\mathbf{v}_1) = \tau(\mathbf{v}_2)$, all $\mathbf{u} \in \mathcal{D}_{\mathbf{U}_L}$, and all $\mathbf{C}_i \in \mathbb{C}$,

$$\begin{aligned} & \tau \left(\left(f_V^L(\mathbf{p}\mathbf{a}_V^{(1)}, \mathbf{u}_V) : V \in \mathbf{C}_i \right) \right) \\ &= \tau \left(\left(f_V^L(\mathbf{p}\mathbf{a}_V^{(2)}, \mathbf{u}_V) : V \in \mathbf{C}_i \right) \right), \end{aligned}$$

where $\mathbf{pa}_V^{(1)}$ and $\mathbf{pa}_V^{(2)}$ are the values corresponding to \mathbf{v}_1 and \mathbf{v}_2 respectively. If \mathbb{D}' is a finer clustering than \mathbb{D} , then that means that $\tau(\mathbf{v}_1) = \tau(\mathbf{v}_2)$ implies $\tau'(\mathbf{v}_1) = \tau'(\mathbf{v}_2)$ for all values $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{D}_{\mathbf{V}_L}$, implying that the above must also hold for τ' . \square

This property implies that the constraints of \mathbb{D} are one-sided, and although finding the most coarse set of clusters may be impossible, any finer set will also work. In the worst case, choosing \mathbb{D} such that $\mathbb{D}_{\mathbf{C}_i} = \mathcal{D}_{\mathbf{C}_i}$ (i.e. every value in their own cluster) would still result in a valid abstraction, as shown below.

Proposition 5. *Consider a low level SCM \mathcal{M}_L and constructive abstraction function τ w.r.t. clusters \mathbb{C} and \mathbb{D} . \mathcal{M}_L is guaranteed to satisfy the AIC w.r.t. τ if and only if $\mathbb{D}_{\mathbf{C}_i} = \{\{\mathbf{c}_i\} : \mathbf{c}_i \in \mathcal{D}_{\mathbf{C}_i}\}$ for all $\mathbf{C}_i \in \mathbb{C}$. \blacksquare*

Proof. The first claim is directly implied by Lemma 6, since this choice of \mathbb{D} simply clusters each value to its own cluster, resulting in a finer clustering than any other clustering.

Without any additional information about \mathcal{M}_L , it is possible for any other choice of clustering to result in \mathcal{M}_L failing to satisfy the AIC. This can be shown by constructing an adversarial example of \mathcal{M}_L for any other choice of intravariabile clustering \mathbb{D}' . Since $\mathbb{D}' \neq \mathbb{D}$, this implies that there exists at least one pair of $\mathbf{c}_1, \mathbf{c}_2 \in \mathbf{C}_i$ for some $\mathbf{C}_i \in \mathbb{C}$ such that \mathbf{c}_1 and \mathbf{c}_2 are in the same cluster $\mathcal{D}_{\mathbf{C}_i}^{(j)} \in \mathbb{D}_{\mathbf{C}_i}$ for some j .

Let $\mathbf{C}_k \in \mathbb{C}$ be a cluster such that there exists at least one pair $V_1 \in \mathbf{C}_i, V_2 \in \mathbf{C}_k$ such that $V_1 \in \mathbf{Pa}_{V_2}$. Since there are no restrictions on \mathcal{M}_L aside from basic assumptions like recursiveness, we can construct one such that such a \mathbf{C}_k exists. Consider the set of functions $\mathcal{F}_{\mathbf{C}_k} = \{f_V^L : V \in \mathbf{C}_k\} \subseteq \mathcal{F}_L$, from \mathcal{M}_L , and denote $\mathcal{F}_{\mathbf{C}_k}(\mathbf{c}_k, \mathbf{u}_L) = (f_V^L(\mathbf{c}_k, \mathbf{u}_L) : V \in \mathbf{C}_k)$. Provided that the domains of \mathbf{V}_H are nontrivial (each variable can take at least two values), there must exist $v_1, v_2 \in \mathcal{D}_{V_k}$, where $V_k = \tau_{\mathbf{C}_k}(\mathbf{C}_k)$, such that $v_1 \neq v_2$, and there exists $\mathbf{c}_{k,1}, \mathbf{c}_{k,2} \in \mathcal{D}_{\mathbf{C}_k}$ such that $\tau_{\mathbf{C}_k}(\mathbf{c}_{k,1}) = v_1$ and $\tau_{\mathbf{C}_k}(\mathbf{c}_{k,2}) = v_2$. Hence, we can construct each function of $\mathcal{F}_{\mathbf{C}_k}$ such that for some setting of $\mathbf{U}_L = \mathbf{u}_L$, $\tau_{\mathbf{C}_k}(\mathcal{F}_{\mathbf{C}_k}(\mathbf{c}_1, \mathbf{u}_L)) \neq \tau_{\mathbf{C}_k}(\mathcal{F}_{\mathbf{C}_k}(\mathbf{c}_2, \mathbf{u}_L))$, violating the AIC. \square

B Background on Causal Abstractions

In this section, we discuss some of the prior works in causal abstractions (Rubenstein et al. 2017; Beckers and Halpern 2019; Beckers, Eberhardt, and Halpern 2019; Geiger, Potts, and Icard 2023; Massidda et al. 2023). In many established causal inference tasks, it is typically assumed that there is a well-specified and known set of variables of interest \mathbf{V} , and nature is modeled by a collection of mechanisms that assign values to each of these variables. However, the definition of \mathbf{V} may not always be clear in practice. In particular, the variables of interest may not align with the features of the data. For example, in an economic system, perhaps data on each individual consumer is collected, but the variable of interest is an aggregate measure like gross domestic product (GDP). In image data, perhaps the pixel values are collected, but the variables of causal interest are related to the objects of the image, not the individual pixels.

Acknowledging that the data is not always provided in the best choice of granularity, existing works of causal abstractions typically define two sets of variables, \mathbf{V}_L and \mathbf{V}_H , which describe the lower level and higher level settings, respectively. For example, \mathbf{V}_L might describe the pixels of an image, while \mathbf{V}_H might describe its structural content. They are typically modeled by corresponding causal models \mathcal{M}_L and \mathcal{M}_H , respectively. In this section, we describe some relevant works in this context and will employ our notation for consistency purposes when their notation differs.

The connection between \mathbf{V}_H and \mathbf{V}_L can be described through a mapping, $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$, between their domains. However, even if τ is known, it is not guaranteed that a model over \mathbf{V}_H , \mathcal{M}_H , is an abstraction of a model over \mathbf{V}_L , \mathcal{M}_L . In short, while τ connects the domains of the variables, there is nothing guaranteeing any kind of connection between the models \mathcal{M}_H and \mathcal{M}_L , be it the functions, the exogenous noise, or the induced distributions.

One of the earliest works that formally discuss abstractions in the context of causal models is Rubenstein et al. (2017), which establishes the idea of *exact transformations*, where a high-level SCM \mathcal{M}_H could be considered an “abstraction” of a low-level SCM \mathcal{M}_L if \mathcal{M}_H is an exact τ -transformation of \mathcal{M}_L . In addition to τ , which connects the domains of \mathbf{V}_H and \mathbf{V}_L , exact transformations connect the two models \mathcal{M}_L and \mathcal{M}_H through their induced interventional distributions. This requires mapping the set of low level interventions \mathcal{I}_L (over \mathbf{V}_L) to their corresponding high-level counterparts \mathcal{I}_H (over \mathbf{V}_H), which is done through another function $\omega : \mathcal{I}_L \rightarrow \mathcal{I}_H$. This idea leads to the following definition.

Definition 20 (Exact Transformation (Rubenstein et al. 2017, Def. 3)). Let \mathcal{M}_L and \mathcal{M}_H be SCMs and $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ be a function. We say that $(\mathcal{M}_H, \mathcal{I}_H)$ is an exact τ -transformation of $(\mathcal{M}_L, \mathcal{I}_L)$ if there exists a surjective order preserving map $\omega : \mathcal{I}_L \rightarrow \mathcal{I}_H$ such that

$$P(\tau(\mathbf{V}_L[\mathbf{x}_L = \mathbf{x}_L])) = P(\mathbf{V}_H[\omega(\mathbf{x}_L = \mathbf{x}_L)]). \quad (61)$$

In this definition, the variables and corresponding distributions of \mathcal{M}_L are linked to those of \mathcal{M}_H through the function τ , and corresponding causal interventions are linked through the function ω .

The interventional sets \mathcal{I}_L and \mathcal{I}_H are called the “allowed” interventions of \mathcal{M}_L and \mathcal{M}_H , respectively. They can be specified to contain any possible intervention and exclude others. Since Eq. 61 only applies in cases where the intervention $\mathbf{X}_L = \mathbf{x}_L$ is contained in \mathcal{I}_L , any intervention that is not in \mathcal{I}_L or \mathcal{I}_H is deemed irrelevant in the context of exact transformations, and no restrictions are placed on their corresponding interventional distributions according to the definition. In an extreme case, if \mathcal{I}_L and \mathcal{I}_H only contained the empty intervention (i.e. $\mathcal{I}_L = \mathcal{I}_H = \{\emptyset\}$), then Eq. 61 would only require that $P(\tau(\mathbf{V}_L)) = P(\mathbf{V}_H)$ and makes no statements about any interventional distributions from \mathcal{L}_2 . As opposed to requiring \mathcal{I}_L and \mathcal{I}_H to contain all interventions, this flexibility allows one to specify which interventions are well-defined, which is important since not every intervention may translate well across an abstraction. Consider the following example for concreteness.

Example 11. Suppose a two-branch government is voting on a law, where Y is a binary variable denoting whether the law is enacted, and X_1 and X_2 are the binary variables representing the votes of the two branches. In this case, $\mathbf{V}_L = \{X_1, X_2, Y\}$. The law is only considered if both branches vote “yes”, so instead of representing the two branches’ votes separately, one could introduce a new variable

$$X_H \leftarrow (X_1 = \text{“yes”}) \wedge (X_2 = \text{“yes”}) \quad (62)$$

as an abstraction of X_1 and X_2 , with $\mathbf{V}_H = \{X_H, Y\}$.

In this case, the low level intervention ($X_1 \leftarrow \text{“yes”}, X_2 \leftarrow \text{“yes”}$) (simultaneously intervening on both X_1 and X_2), would map to a high level intervention ($X_H \leftarrow 1$). Eq. 61 then dictates that

$$P(\tau(\mathbf{V}_{L[X_1=\text{“yes”}, X_2=\text{“yes”}]})) = P(\mathbf{V}_{H[X_H=1]}). \quad (63)$$

However, an intervention like ($X_1 \leftarrow \text{“yes”}$), which only intervenes on X_1 , does not have a corresponding high level counterpart. The value of X_H under this intervention would still depend on X_2 . Hence, the intervention ($X_1 \leftarrow \text{“yes”}$) should be excluded from \mathcal{I}_L , which implies no restrictions on $P(\tau(\mathbf{V}_{L[X_1=\text{“yes”}]}))$. ■

For any interventional set \mathcal{I} , there exists a natural partial ordering \leq such that $i \leq j$ for $i, j \in \mathcal{I}$ if and only if the interventional values of i are a subset of those in j (e.g., $(A \leftarrow a, B \leftarrow b) \leq (A \leftarrow a, B \leftarrow b, C \leftarrow c)$). Given the orderings \leq_L and \leq_H of \mathcal{I}_L and \mathcal{I}_H , respectively, the order preserving property of ω is defined to mean that $i \leq_L j$ for $i, j \in \mathcal{I}_L$ implies $\omega(i) \leq_H \omega(j)$. This property enforces a kind of regularity condition on ω ensuring that low-level interventions are still related when translated to the higher level. For example, an intervention of $(A \leftarrow a, B \leftarrow b, C \leftarrow c)$ on the low level may map through ω to an intervention $(X \leftarrow x, Z \leftarrow z)$ on the high level. If we consider the same intervention, but with one more added value, such as $(A \leftarrow a, B \leftarrow b, C \leftarrow c, D \leftarrow d)$, we would expect that ω would map it to a similar intervention, possible with more values on the high level, such as $(X \leftarrow x, Z \leftarrow z, Y \leftarrow y)$. The order preserving property of ω prevents it from mapping the intervention to one with fewer values such as $(X \leftarrow x)$ or ones with different values like $(X \leftarrow x, Z \leftarrow z')$.

With all of these properties, exact τ -transformations establish an important foundational property expected from all abstractions, namely, that the abstraction mapping τ commutes with applied interventions (Eq. 61), as illustrated in Fig. 14.

Example 12. Consider a low level SCM $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle$ that models an alarm system. Suppose $\mathbf{V}_L = \{E, S, A\}$, all binary, where the alarm rings ($A = 1$) if either there is an earthquake ($E = 1$) or smoke from a fire ($S = 1$), with some possible noise. Formally, the causal mechanisms are described as follows:

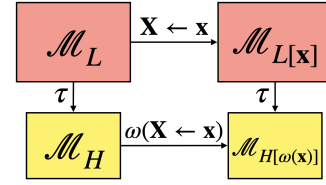


Figure 14: Illustration of interventional commutativity. Applying low-level intervention $\mathbf{X} \leftarrow \mathbf{x}$ followed by abstraction τ is equivalent to first applying τ , followed by applying high-level intervention $\omega(\mathbf{X} \leftarrow \mathbf{x})$.

$$\mathbf{U}_L = \{U_E, U_S, U_A\} \quad (64)$$

$$\mathbf{V}_L = \{E, S, A\} \quad (65)$$

$$\mathcal{F}_L = \begin{cases} E \leftarrow f_E^L(u_E) & = u_E \\ S \leftarrow f_S^L(u_S) & = u_S \\ A \leftarrow f_A^L(e, s, u_A) & = (e \vee s) \oplus u_A \end{cases} \quad (66)$$

$$P(\mathbf{U}_L) : P(U_E = 1) = P(U_S = 1) = P(U_A = 1) = 0.5 \quad (67)$$

Now suppose instead of considering both earthquake and fire individually, we would like to abstract both of these events into a less granular variable D , representing whether or not some disaster has occurred. That is, $D = E \vee S$. To be precise, this means the high level variables can be defined as $\mathbf{V}_H = \{D, A_H\}$, and τ can be defined such that

$$(D, A_H) \leftarrow \tau(e, s, a) = (e \vee s, a). \quad (68)$$

Now consider the following SCM $\mathcal{M}_H^{(1)}$ defined over these variables:

$$\mathbf{U}_1 = \{U_D, U_A\} \quad (69)$$

$$\mathbf{V}_H = \{D, A_H\} \quad (70)$$

$$\mathcal{F}_1 = \begin{cases} D \leftarrow f_D^1(u_D) & = u_D \\ A_H \leftarrow f_{A_H}^1(d, u_A) & = d \oplus u_A \end{cases} \quad (71)$$

$$P(\mathbf{U}_1) : P(U_D = 1) = 0.75, P(U_A = 1) = 0.5 \quad (72)$$

Further suppose that the lists of allowed interventions are

$$\begin{aligned} \mathcal{I}_L = \{ & \emptyset, (E \leftarrow 1), (S \leftarrow 1), (E \leftarrow e, S \leftarrow s), \\ & (A \leftarrow a), (E \leftarrow 1, A \leftarrow a), \\ & (S \leftarrow 1, A \leftarrow a), (E \leftarrow e, S \leftarrow s, A \leftarrow a) \}, \end{aligned} \quad (73)$$

and

$$\mathcal{I}_H^{(1)} = \{ \emptyset, (D \leftarrow d), (A_H \leftarrow a), (D \leftarrow d, A_H \leftarrow a) \}, \quad (74)$$

for all settings of e, s, a, d . The partial ordering of these interventions are in the listed order.

One can verify that $(\mathcal{M}_H^{(1)}, \mathcal{I}_H^{(1)})$ is an exact τ -transformation of $(\mathcal{M}_L, \mathcal{I}_L)$. For ω , we can choose one that maps any combination of $E \leftarrow e, S \leftarrow s, A \leftarrow a$ to $D \leftarrow e \vee s, A_H \leftarrow a$. Any case with $E \leftarrow 1$ or $S \leftarrow 1$ automatically maps to $D \leftarrow 1$. We note that it is order preserving (e.g., $(E \leftarrow 0, S \leftarrow 0) \leq_L (E \leftarrow 0, S \leftarrow 0, A \leftarrow 0)$ and

$(D \leftarrow 0) \leq_H (D \leftarrow 0, A_H \leftarrow 0)$), and one can verify that Eq. 61 holds. For example,

$$P(\tau(E = 1, S = 0)) = 0.75 = P(D = 1) \quad (75)$$

and

$$P(\tau(A_{E=1, S=0} = 1)) = 0.5 = P(A_{H[D=1]} = 1). \quad (76)$$

One notable property that allows this to occur is that $E \leftarrow 0$ and $S \leftarrow 0$ are not valid interventions according to \mathcal{I}_L , since the corresponding intervention mapped by ω is ambiguous as discussed earlier in Ex. 11.

Now consider an alternative model $\mathcal{M}_H^{(2)}$, described as follows:

$$\mathbf{U}_2 = \{U_D, U_A\} \quad (77)$$

$$\mathbf{V}_H = \{D, A_H\} \quad (78)$$

$$\mathcal{F}_2 = \begin{cases} D \leftarrow f_D^2(a, u_E) & = u_D \\ A_H \leftarrow f_{A_H}^2(u_A) & = u_A \end{cases} \quad (79)$$

$$P(\mathbf{U}_2) : P(U_D = 1) = 0.75, P(U_A = 1) = 0.5. \quad (80)$$

It turns out that, if $\mathcal{I}_H^{(2)} = \mathcal{I}_H^{(1)}$, then $(\mathcal{M}_H^{(2)}, \mathcal{I}_H^{(2)})$ is an exact τ -transformation of $(\mathcal{M}_L, \mathcal{I}_L)$ for the same choice of τ , even though in $\mathcal{M}_H^{(2)}$, the causal relationship between D and A_H no longer exists. In fact, consider another case where $\mathcal{I}_L = \mathcal{I}_H^{(2)} = \{\emptyset\}$, that is, only the empty intervention is allowed on either level, and $\omega(\emptyset) = \emptyset$. Perhaps surprisingly, in this case, $(\mathcal{M}_H^{(2)}, \mathcal{I}_H^{(2)})$ is an exact τ -transformation of $(\mathcal{M}_L, \mathcal{I}_L)$ for several other counterintuitive choices of τ as well. For instance, we can choose $(D, A_H) \leftarrow \tau(e, s, a) = (-e \vee a, s)$, which does not even map A to A_H . We can even choose one that is not consistent across variables, for example $\tau(e, s, a)$ can map $(0, 1, 0), (0, 0, 1), (1, 0, 1)$ to $(1, 0); (0, 0, 0), (0, 1, 1), (1, 1, 0)$ to $(1, 1); (1, 1, 1)$ to $(0, 0);$ and $(1, 0, 0)$ to $(0, 1)$. One can verify that, in both cases, this still results in an exact τ -transformation. Furthermore, one can imagine changing the names of the variables to describe something arbitrarily different. In fact, it seems that $\mathcal{M}_H^{(2)}$ and \mathcal{M}_L are completely unrelated. ■

The heart of the issue raised in the previous example is that in cases where several choices of \mathbf{v}_L have the same probability, one can still obtain a valid abstraction without violating Eq. 61 by rearranging values mapped by τ that have the same probability, even if the resulting rearrangement has no causal interpretation. Further when the allowed interventions are sparse, Eq. 61 is required to hold on fewer distributions, resulting in a weaker connection between the high and low-level models. In these cases, the definition of exact τ -transformations becomes weak and can often no longer be used to define abstractions in any intuitive sense.

Building on the work of Rubenstein et al. (2017), Beckers and Halpern (2019) introduced several refined definitions that resolved these issues, including the notion of τ -abstractions. We rewrite the definitions as shown below.

Definition 14 (Beckers and Halpern 2019, Def. 3.12)). Some relevant notation is defined as follows:

- Given a set of variables $\mathbf{V}, \mathbf{X} \subseteq \mathbf{V}$, and $\mathbf{x} \in \mathcal{D}_{\mathbf{X}}$, let $\text{Rst}(\mathbf{V}, \mathbf{x}) = \{\mathbf{v} \in \mathcal{D}_{\mathbf{V}} : \mathbf{v} \text{ is consistent with } \mathbf{x}\}$.
- Given variables \mathbf{V}_L and \mathbf{V}_H , mapping $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$, and value set $\mathbf{T} \subseteq \mathcal{D}_{\mathbf{V}_L}$, denote $\tau(\mathbf{T}) = \{\tau(\mathbf{v}_L) : \mathbf{v}_L \in \mathbf{T}\}$.
- Given allowed interventions \mathcal{I}_L and \mathcal{I}_H over \mathbf{V}_L and \mathbf{V}_H respectively, define $\omega_\tau : \mathcal{I}_L \rightarrow \mathcal{I}_H$ such that $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L) = \mathbf{X}_H \leftarrow \mathbf{x}_H$, where $\tau(\text{Rst}(\mathbf{V}_L, \mathbf{x}_L)) = \text{Rst}(\mathbf{V}_H, \mathbf{x}_H)$.

Definition 15 (τ -Abstraction (Beckers and Halpern 2019, Def. 3.13)). Let $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle$ and $\mathcal{M}_H = \langle \mathbf{U}_H, \mathbf{V}_H, \mathcal{F}_H, P(\mathbf{U}_H) \rangle$ be two SCMs. Let \mathcal{I}_L and \mathcal{I}_H be the sets of allowed interventions respectively. Given $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$, we say that $(\mathcal{M}_H, \mathcal{I}_H)$ is a τ -abstraction of $(\mathcal{M}_L, \mathcal{I}_L)$ if:

1. τ is surjective;
2. There exists surjective $\tau_{\mathbf{U}} : \mathcal{D}_{\mathbf{U}_L} \rightarrow \mathcal{D}_{\mathbf{U}_H}$ that is compatible with τ , i.e.

$$\tau(\mathcal{M}_{L[\mathbf{X}_L \leftarrow \mathbf{x}_L]}(\mathbf{u}_L)) = \mathcal{M}_{H[\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L)]}(\tau_{\mathbf{U}}(\mathbf{u}_L)), \quad (59)$$

for all $\mathbf{u}_L \in \mathcal{D}_{\mathbf{U}_L}$ and all $(\mathbf{X}_L \leftarrow \mathbf{x}_L) \in \mathcal{I}_L$;

3. $\mathcal{I}_H = \omega_\tau(\mathcal{I}_L)$. ■

Definition 16 (Strong τ -Abstraction (Beckers and Halpern 2019, Def. 3.15)). We say that \mathcal{M}_H is a strong τ -abstraction of \mathcal{M}_L if $(\mathcal{M}_H, \mathcal{I}_H)$ is a τ -abstraction of $(\mathcal{M}_L, \mathcal{I}_L)$ and $\mathcal{I}_H = \mathcal{I}_H^*$. ■

We discuss each point of Def. 15 in detail:

1. The surjectivity of τ does not add any mathematical benefits but is required as a property because \mathbf{V}_H is expected to be less “complex” than \mathbf{V}_L if \mathcal{M}_H is to be called an abstraction of \mathcal{M}_L .
2. The addition of $\tau_{\mathbf{U}}$ is the major constraint added to τ -abstractions when compared to exact transformations. By establishing a connection between the lower and higher level exogenous variables, $\tau_{\mathbf{U}}$ ensures that the distributions of \mathbf{V}_H meaningfully correspond to those of \mathbf{V}_L . This, in fact, fixes the issue described earlier in Ex. 12, since even if Eq. 61 holds, it is not necessarily the case that Eq. 59 will hold unless the values of \mathbf{U}_L that are used to compute the l.h.s. of Eq. 61 match those of $\tau_{\mathbf{U}}(\mathbf{U}_L)$ that are used to compute the r.h.s. This point is illustrated in Ex. 13 below.
3. With \mathcal{I}_L and ω_τ fixed, \mathcal{I}_H should be fixed to $\omega_\tau(\mathcal{I}_H)$ to remain consistent.

Def. 16 further fixes the issue of \mathcal{I}_L containing too few interventions. When \mathcal{I}_H is maximal, then every submodel $\mathcal{M}_{H[\mathbf{x}_H]}$ has a corresponding submodel $\mathcal{M}_{L[\mathbf{x}_L]}$ such that $\omega_\tau(\mathbf{X}_L \leftarrow \mathbf{x}_L) = \mathbf{X}_H \leftarrow \mathbf{x}_H$, as described in Eq. 59.

Example 13 (Example 12 continued). Consider the high level models $\mathcal{M}_H^{(1)}$ and $\mathcal{M}_H^{(2)}$ from Ex. 12. Using the definition of $\tau(e, s, a) = (e \vee s, a)$, ω_τ is forced to take the mapping as specified earlier where

$$E \leftarrow e, S \leftarrow s, A \leftarrow a \quad (81)$$

maps to

$$D \leftarrow e \vee s, A_H \leftarrow a. \quad (82)$$

Note that in this case, $(\mathcal{M}_H^{(1)}, \mathcal{I}_H^{(1)})$ is still a τ -abstraction of $(\mathcal{M}_L, \mathcal{I}_L)$. Specifically, choose τ_U such that

$$(U_D, U_A) \leftarrow \tau_U(u_E, u_S, u_A) = (u_E \vee u_S, u_A). \quad (83)$$

One can verify, for example, that

$$\tau(\mathcal{M}_{L[E=e, S=s]}(\mathbf{u}_L)) = \mathcal{M}_{H[D=e \vee s]}^{(1)}(\tau_U(\mathbf{u}_L)) \quad (84)$$

for all values of e, s , and \mathbf{u}_L , aligning with Eq. 59.

On the other hand, $(\mathcal{M}_H^{(2)}, \mathcal{I}_H^{(2)})$ is not a τ -abstraction of $(\mathcal{M}_L, \mathcal{I}_L)$. For example, Eq. 59 states that

$$\tau(\mathcal{M}_{L[E=0, S=0]}(\mathbf{u}_L)) = \mathcal{M}_{H[D=0]}^{(2)}(\tau_U(\mathbf{u}_L)) \quad (85)$$

and

$$\tau(\mathcal{M}_{L[E=1, S=0]}(\mathbf{u}_L)) = \mathcal{M}_{H[D=1]}^{(2)}(\tau_U(\mathbf{u}_L)). \quad (86)$$

However, note that

$$\tau(A_{E=0, S=0}(u_A)) \neq \tau(A_{E=1, S=0}(u_A)). \quad (87)$$

For example, when $E = 0, S = 0, U_A = 0$, then f_A^L will assign $A = 0$, but when $E = 1, S = 0, U_A = 0$, f_A^L will assign $A = 1$. On the other hand,

$$A_{H[D=0]}^{(2)}(\tau_U(u_A)) = A_{H[D=1]}^{(2)}(\tau_U(u_A)) \quad (88)$$

for any choice of τ_U because $f_{A_H}^2$ does not take D as an input. This contradicts the equalities enforced by Eqs. 85 and 86. ■

The definitions introduced so far are effective at describing one SCM as an abstraction of another. For example, if two SCMs \mathcal{M}_H and \mathcal{M}_L are provided, as well as the function τ , Def. 15 can be used to decide whether \mathcal{M}_H is indeed an abstraction of \mathcal{M}_L . However, this may not be particularly useful in cases where the higher level model \mathcal{M}_H is not known in advance, and one would like to find or learn such an abstraction. Beckers and Halpern (2019) makes an important step in the direction of applying such works by defining a more concrete class of abstractions that can be obtained by construction.

Definition 17 (Constructive τ -Abstraction (Beckers and Halpern 2019, Def. 3.19)). \mathcal{M}_H is a constructive τ -abstraction of \mathcal{M}_L if \mathcal{M}_H is a strong τ -abstraction of \mathcal{M}_L , and there exists a partition of $\mathbf{V}_L, \mathbb{C} = \{\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_{n+1}\}$ (where $n = |\mathbf{V}_H|$) with nonempty \mathbf{C}_1 to \mathbf{C}_n , such that τ can be decomposed as $\tau = (\tau_{\mathbf{C}_1}, \tau_{\mathbf{C}_2}, \dots, \tau_{\mathbf{C}_n})$, where each $\tau_{\mathbf{C}_i} : \mathcal{D}_{\mathbf{C}_i} \rightarrow \mathcal{D}_{V_{H,i}}$ maps the i th partition to the i th variable of \mathbf{V}_H . ■

In this definition, variables of \mathbf{V}_L are specifically partitioned into clusters $\mathbf{C}_1, \dots, \mathbf{C}_{n+1}$, and τ is defined such that each cluster \mathbf{C}_i maps to a high level variable $V_i \in \mathbf{V}_H$. The definition of τ and corresponding high level space \mathbf{V}_H are concretely defined in this definition. We leverage a similar concept in this paper, allowing the higher level variables \mathbf{V}_H (and correspondingly, τ), to be defined by construction based on predetermined clusters of lower level variables.

B.1 Comparisons with Sec. 2

The approach used in this work leverages similar ideas to constructive τ -abstractions (Def. 17) for the purpose of obtaining the high level model \mathcal{M}_H constructively. Notably, the intervariable clusters in Def. 5 partition the variable space in the same way, and the corresponding choice of τ from Def. 6 is defined around these clusters, by utilizing a different subfunction $\tau_{\mathbf{C}_i}$ for each intervariable cluster \mathbf{C}_i , similar to Def. 17. For these reasons, any choice of τ that follows Def. 6 is called a constructive abstraction function.

Still, the major difference is that Def. 17 focuses on the relationship between the full models \mathcal{M}_L and \mathcal{M}_H , while Def. 6 only defines the mapping τ that connects the variable spaces \mathbf{V}_L and \mathbf{V}_H . That is, Def. 6 by itself makes no claims about how other aspects of \mathcal{M}_L (such as the functions \mathcal{F}_L or exogenous noise $P(\mathbf{U}_L)$) relate to \mathcal{M}_H , other than the variables \mathbf{V}_L and \mathbf{V}_H . This is a new approach to abstraction work. Note that exact transformations (Def. 20) and τ -abstractions (Def. 15) place no requirements on the definition of τ , and Def. 17 only requires that τ can be decomposed relative to a partition. Indeed, placing requirements on τ reduces its generality, but ensuring that τ follows the form illustrated in Def. 6 has several advantages:

1. [**Query-Specific Abstractions**] The primary purpose of Def. 6 is to introduce a relaxed notion of abstractions that are defined on specific distributions of the PCH. As opposed to exact transformations and τ -abstractions (including constructive ones), which focus on the entire SCMs \mathcal{M}_L and \mathcal{M}_H , the concept of Q - τ consistency (Def. 7) allows one to define “partial” abstractions. For instance, a choice of \mathcal{M}_H can be considered an abstraction of \mathcal{M}_L for Q_1 but not Q_2 if \mathcal{M}_H is Q_1 - τ consistent with \mathcal{M}_L but not Q_2 - τ consistent. This subtlety is lost in τ -abstractions for example, where any mismatch of Eq. 59 disqualifies \mathcal{M}_H from being considered an abstraction of \mathcal{M}_L . See Example 14 below for a more concrete explanation on this distinction. Indeed, when \mathcal{M}_H is \mathcal{L}_3 - τ consistent with \mathcal{M}_L , that is, \mathcal{M}_H is Q - τ consistent with \mathcal{M}_L on every possible counterfactual query, then it turns out that \mathcal{M}_H behaves like a constructive- τ abstraction of \mathcal{M}_L (see Prop. 1).¹⁷

Defining abstractions on the level of individual queries enables a more practical approach to learning abstractions. The true model \mathcal{M}_L is rarely available in practice, and instead, one is often given data from \mathcal{M}_L from its induced distributions (e.g., the observational distribution $P(\mathbf{V}_L)$). As discussed in Sec. 3, one would ideally be able to construct a high-level model \mathcal{M}_H that is Q - τ consistent with the available distributions. Such a model may not be Q - τ consistent with \mathcal{M}_L on choices of Q that were not provided in the data, but through Alg. 2, one can determine precisely which choices of Q do indeed match across abstractions. \mathcal{M}_H can be considered an abstraction of \mathcal{M}_L for those cases.

¹⁷Note that all τ -abstractions are exact τ -transformations, a result from Beckers and Halpern (2019). However, exact transformations are not necessarily \mathcal{L}_3 - τ consistent because Eq. 61 is focused on \mathcal{L}_2 and is oblivious to the counterfactual level.

2. **[Natural Interventional Mapping]** As opposed to exact transformations and τ -abstractions, the concept of a function ω that maps between intervention spaces \mathcal{I}_L and \mathcal{I}_H is no longer required since the corresponding high-level intervention of a low-level one is straightforward. Specifically, the intervention $\mathbf{X}_L \leftarrow \mathbf{x}_L$ maps to $\tau(\mathbf{X}_L) \leftarrow \tau(\mathbf{x}_L)$ (see Lem. 3). Additionally, the concept of “allowed interventions” is no longer needed. One can simply set \mathcal{I}_H as the set of all high level interventions \mathcal{I}_H^* , and choose \mathcal{I}_L as the set of corresponding interventions that map to \mathcal{I}_H^* (interventions of unions of clusters, as shown in Lem. 4). For these reasons, the presentation in Sec. 2 does not include any references of \mathcal{I}_L , \mathcal{I}_H or ω , leaving the focus of the discussion on the abstraction of the variables \mathbf{V}_L .
3. **[True Constructiveness]** Despite the progress that constructive τ -abstractions (Def. 17) make in the direction of constructively building \mathcal{M}_H from \mathcal{M}_L , the definition alone does not accomplish this task. By defining τ as a mapping across clusters of both variables and values, one can leverage Alg. 1 from this paper to obtain the high level model \mathcal{M}_H when given the low level \mathcal{M}_L . When \mathcal{M}_L is not provided, and data from \mathcal{M}_L is provided instead, one can use Alg. 2 to obtain a model \mathcal{M}_H which is still an abstraction of \mathcal{M}_L on identifiable queries. This approach is implementable in practice leveraging neural optimization approaches, and the experiments provided in Sec. 5 demonstrate their applicability.
4. **[Intuitive Abstractions]** Abstractions that are constructed with a choice of τ that does not follow Def. 6 can be quite esoteric (see Ex. 20 in App. E). The concept of clustering is intuitive, and the relationship between \mathbf{V}_H and \mathbf{V}_L when τ is a constructive abstraction function is straightforward and interpretable. Intervariable clusters can be determined based on the needs of the task or constructed algorithmically via Alg. 3 (in App. D.1). These clusters also have a natural connection with cluster causal diagrams (Anand et al. 2023), as illustrated in Sec. 3. Intravariabile clusters are strongly tied to invariances in the data (two low-level values that are clustered together will map to the same high-level value), leading to a strong connection with representation learning (see Sec. 4 and App. D.3).

Consider the following example comparing constructive τ -abstractions to the concept of Q - τ consistency.

Example 14 (Example 12 continued). Recall from Example 12 that the SCM $\mathcal{M}_H^{(1)}$ as described by Eqs. 69 to 72 is a τ -abstraction (and therefore also an exact τ -transformation) of \mathcal{M}_L described by Eqs. 64 to 67.

It turns out that this choice of τ , defined such that

$$(D, A_H) \leftarrow \tau(e, s, a) = (e \vee s, a), \quad (89)$$

is actually a constructive abstraction function with the intervariable clusters

$$\mathbb{C} = \{\mathbf{C}_1 = \{E, S\}, \mathbf{C}_2 = \{A\}\} \quad (90)$$

and intravariabile clusters

$$\mathbb{D} = \{\mathbb{D}_{\mathbf{C}_1}, \mathbb{D}_{\mathbf{C}_2}\}, \quad (91)$$

where

$$\mathbb{D}_{\mathbf{C}_1} = \begin{cases} d_0 & = \{(E = 0, S = 0)\}, \\ d_1 & = \{(E = 0, S = 1), (E = 1, S = 0), \\ & (E = 1, S = 1)\}, \end{cases} \quad (92)$$

and $\mathbb{D}_{\mathbf{C}_2}$ retains the same values of \mathcal{D}_A .

It is then easy to verify that $\mathcal{M}_H^{(1)}$ is also \mathcal{L}_3 - τ consistent with \mathcal{M}_L . For example, $P^{\mathcal{M}_L}(A_{E=0, S=0} = 1 \mid A = 1, E = 0, S = 1) = 0$, which is the counterfactual probability that the alarm would ring had neither earthquake nor fire occurred, given the reality that the alarm indeed rang when there was a fire but no earthquake. It is also true that the equivalent query mapped across τ (from Def. 7) is consistent, that is, $P^{\mathcal{M}_H^{(1)}}(A_{H[D=0]} = 1 \mid A = 1, D = 1) = 0$. This \mathcal{L}_3 - τ consistency is a consequence of Prop. 1.

Now consider another high level model $\mathcal{M}_H^{(3)}$ also defined over \mathbf{V}_H from the same τ .

$$\mathbf{U}_3 = \{U_D, U_{A_0}, U_{A_1}\} \quad (93)$$

$$\mathbf{V}_H = \{D, A_H\} \quad (94)$$

$$\mathcal{F}_3 = \begin{cases} f_D^3(u_E) & = u_D \\ f_{A_H}^3(d, u_{A_0}, u_{A_1}) & = \begin{cases} u_{A_0} & d = 0 \\ u_{A_1} & d = 1 \end{cases} \end{cases} \quad (95)$$

$$P(\mathbf{U}_3) : P(U_D = 3) = 0.75, \quad (96)$$

$$P(U_{A_0} = 1) = P(U_{A_1} = 1) = 0.5 \quad (97)$$

Note that $P^{\mathcal{M}_H^{(3)}}(A_{H[D=0]} = 1 \mid A = 1, D = 1) = P(U_{A_0} = 1) = 0.5$, which is inconsistent with the result from \mathcal{M}_L . Hence, $\mathcal{M}_H^{(3)}$ is not \mathcal{L}_3 - τ consistent with \mathcal{M}_L , nor is it a τ -abstraction. Still, careful analysis of $\mathcal{M}_H^{(3)}$ reveals that it is still \mathcal{L}_2 - τ consistent with \mathcal{M}_L . For example, $P^{\mathcal{M}_H^{(3)}}(A_{H[D=1]} = 1) = 0.5 = P^{\mathcal{M}_L}(A_{E=1, S=0} = 1)$. Therefore, $\mathcal{M}_H^{(3)}$ may still be a valid abstraction of \mathcal{M}_L if used to infer layer 2 or interventional quantities. This highlights the limitations of a definition of abstractions that works on the level of the SCM, such as τ -abstractions. If the user of the model is only interested in interventional quantities, it may be premature to discount $\mathcal{M}_H^{(3)}$ as an invalid abstraction. ■

B.2 Learning Abstraction Functions

The previous section discussed the benefits of defining the abstraction function τ to be a constructive abstraction function from Def. 6, with many of the reasons leading to the ability to learn the higher-level model \mathcal{M}_H . Nonetheless, there are works that solve the inversion version of this problem, namely, where information about \mathcal{M}_H is given but the function τ is unknown.

For example, Zennaro et al. (2023) solves the problem where, given both low and high-level models \mathcal{M}_L and \mathcal{M}_H , as well as intervariable clusters \mathbb{C} , the goal is to learn subfunctions $\hat{\tau}_{\mathbf{C}_i}$ for each $\mathbf{C}_i \in \mathbb{C}$ that satisfies desired properties of abstractions. Notably, the key properties required in this case are the commutativity of interventions, as illustrated in

Fig. 14, as well as surjectivity of the functions. While a blind search over possible choices of τ is intractable, the paper cleverly models the commutativity property as a continuous error term that decreases as commutativity is closer to being achieved. They further model $\hat{\tau}$ as a neural network and train it to minimize the error term using gradient descent, while regularizing the objective to enforce surjectivity. This approach allows for a more robust solution to the issue of intractability of specifying intravariabile clusters, as discussed in Sec. 4, albeit with the additional requirement that data is available from the high-level model \mathcal{M}_H .

More recently, Felekis et al. (2024) solved a similar problem with the requirement of having the full specification of \mathcal{M}_L and \mathcal{M}_H relaxed. The proposed approach leverages principles of optimal transport to acquire the abstraction map $\hat{\tau}$ using only interventional data from the two models. Since \mathcal{M}_L and \mathcal{M}_H are not provided, the optimization procedure instead enforces do-calculus constraints. In this setting, it is desirable to provide as much expressiveness as possible in the modeling of $\hat{\tau}$ (e.g., with universal approximators such as neural networks), and the optimization may output a result for $\hat{\tau}$ that is not necessarily a constructive abstraction function.

C Experimental Details

This section provides details about our experimental setup and models. Our pipeline is primarily built with PyTorch (Paszke et al. 2017), and training is facilitated using PyTorch Lightning (Falcon and Cho 2020).

C.1 Nutrition Experimental Setup

The nutrition experiment in Sec. 5.1 is a toy study of various individuals and their diets, performed over variables $\mathbf{V}_L = \{R, D, C, F, P, B\}$, where R is restaurant attended, D is dish ordered at that restaurant, C is carbohydrates of the dish in grams, F is fat of the dish in grams, P is protein of the dish in grams, and B is the BMI of the individual. R and D have domains of size 32, indicating 32 different options of restaurant and dishes. C , F , P , and B are real valued numbers. The data generating model $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle$ is described below.

$$\mathbf{U}_L = \{U_R, U_D, U_{RB}, U_{N1}, U_{N2}, U_{N3}, U_B\}$$

$$\mathbf{V}_L = \{R, D, C, F, P, B\}$$

$$\mathcal{F}_L = \{$$

$$f_R^L(u_R, u_{RB}) = (u_R + 16 \cdot u_{RB}) \% 32$$

$$f_D^L(r, u_D) = (r + u_D) \% 32$$

$$f_C^L(d, u_{N1}, u_{N2}, u_{N3}) = 216 \cdot u_{N1} [d \% 3]$$

$$\cdot \left(0.25 \left(\left\lfloor \frac{f}{16} \right\rfloor \oplus u_{N3} \right) + 1 \right) + 9 \cdot u_{N2}$$

$$f_P^L(d, u_{N1}, u_{N2}, u_{N3}) = 216 \cdot u_{N1} [(d + 1) \% 3]$$

$$\cdot \left(0.25 \left(\left\lfloor \frac{f}{16} \right\rfloor \oplus u_{N3} \right) + 1 \right) + 9 \cdot u_{N2}$$

$$f_F^L(d, u_{N1}, u_{N2}, u_{N3}) = 96 \cdot u_{N1} [(d + 2) \% 3]$$

$$\cdot \left(0.25 \left(\left\lfloor \frac{f}{16} \right\rfloor \oplus u_{N3} \right) + 1 \right) + 4 \cdot u_{N2}$$

$$f_B^L(c, f, p, u_B, u_{RB}) =$$

$$\left(\left(\frac{c}{9} + \frac{f}{4} + \frac{p}{9} + 3 \cdot u_{RB} \right) - 30 \right) \cdot (-1)^{u_B} + 25$$

$$P(\mathbf{U}_L) = \begin{cases} P(U_R = u_R) \\ = \begin{cases} \frac{3}{64} & u_R \in \{0, 1, \dots, 15\} \\ \frac{1}{64} & u_R \in \{16, 17, \dots, 31\} \\ 0 & \text{otherwise} \end{cases} \\ P(U_D = u_D) \\ = \begin{cases} \frac{1}{7} & u_D \in \{-3, -2, -1, 0, 1, 2, 3\} \\ 0 & \text{otherwise} \end{cases} \\ U_{RB} \sim \text{Bernoulli}(0.25) \\ U_{N1} \sim \text{Dirichlet}(4, 1, 1) \\ U_{N2} \sim \text{Unif}(0, 1) \\ U_{N3} \sim \text{Bernoulli}(0.1) \\ U_B \sim \text{Bernoulli}(0.1) \end{cases}$$

where $\%$ indicates the “modulo” operator, \oplus is the binary XOR operator, and $u_{N1}[i]$ denotes the i th index of u_{N1} , which is a 3-dimensional variable. In the experiments, R and D are formatted as one-hot vectors.

For the abstraction of \mathcal{M}_L , we choose intervariable clusters $\mathbb{C} = \{D_H = \{D\}, Z = \{C, F, P\}, B_H = \{B\}\}$, where R is abstracted away, D and B are put into their own clusters, and C , F , and P are clustered into a new variable Z , called “calories”. Intravariabile clusters \mathbb{D} are chosen such that the values of each intervariable cluster are divided into two sets (i.e. D_H , Z , B_H are all binary variables). Specifically,

$$\tau_{D_H}(d) = \begin{cases} 0 & d \in \{0, 1, \dots, 15\} \\ 1 & d \in \{16, 17, \dots, 31\} \end{cases}$$

$$\tau_Z(c, f, p) = \mathbf{1}\{4c + 9f + 4p \geq 1080\}$$

$$\tau_{B_H}(b) = \mathbf{1}\{b \geq 25\}$$

For example,

$$\begin{aligned} \tau(R = 7, D = 24, C = 80, F = 70, P = 40, B = 32) \\ = (D_H = 1, Z = 1, B_H = 1). \end{aligned}$$

The high level variables $\mathbf{V}_H = \tau(\mathbf{V}_L)$ are defined to be $\{D_H, Z, B_H\}$.

The causal diagram \mathcal{G} over \mathbf{V}_L and the corresponding C-DAG $\mathcal{G}_{\mathbb{C}}$ over \mathbf{V}_H are shown in Fig. 5. The query of interest is $Q = P(B_{F=f} \geq 25)$, where f is any arbitrary unhealthy food option $f \in \{16, 17, \dots, 31\}$. The query can be interpreted as the probability of someone being overweight if they are forced (intervened) to eat unhealthy food. The corresponding query on the higher level is $\tau(Q) = P(B_{H[D_H=1]} = 1)$, computed from Eq. 10 in Def. 7. The task is to identify and estimate the query Q given observational data $P(\mathbf{V}_L)$ and causal diagram \mathcal{G} . We test three approaches:

1. The first approach is to directly identify and estimate Q from $P^{\mathcal{M}_L}(\mathbf{V}_L)$ and \mathcal{G} . The NeuralID algorithm (Xia, Pan, and Bareinboim 2023, Alg. 1) is used on these inputs. In this approach, the \mathcal{G} -NCM \widehat{M} is fitted over

the variables \mathbf{V}_L and graph \mathcal{G} and is trained on the data from $P^{\mathcal{M}_L}(\mathbf{V}_L)$. It is optimized for the identification/estimation tasks, and Q is directly queried from \widehat{M} .

2. The second approach is identical to the first, except values of \mathbf{V}_L (specifically C , F , P , and B) are all normalized between -1 and 1. The query is reformulated to $Q = P(B_{D=d} \geq 0)$, and the NCM \widehat{M} is modified to use these normalized values¹⁸.
3. The third approach is the abstraction approach, which instead identifies and estimates $\tau(Q)$ from $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$ and \mathcal{G}_C , running Alg. 2 on these inputs. That is, the \mathcal{G}_C -NCM \widehat{M} is fitted over variables \mathbf{V}_H and graph \mathcal{G}_C and is trained on the data of $P^{\mathcal{M}_L}(\mathbf{V}_L)$ passed through τ . Note that this approach already requires fewer assumptions by using the C-DAG \mathcal{G}_C instead of the full causal diagram \mathcal{G} . The end result should be theoretically equivalent to identifying and estimating Q due to Corol. 1.

The experimental results are shown in Fig. 10, where the first approach is shown in red, the second in yellow, and the third in blue.

C.2 Nutrition Models and Hyperparameters

All three models used in the nutrition experiment are GAN-NCMs from Xia, Pan, and Bareinboim (2023), which leverage generative adversarial networks (GAN) (Goodfellow et al. 2014). The GAN-NCMs for the first two approaches without abstractions are SCMs $\widehat{M} = \langle \widehat{U}, \mathbf{V}_L, \widehat{\mathcal{F}}, P(\widehat{U}) \rangle$ fitted to the graph \mathcal{G} (Fig. 5 (left)). Each function $f_V \in \widehat{\mathcal{F}}$ is a feedforward neural network with 3 hidden layers of width 32, with layer normalization applied (Ba, Kiros, and Hinton 2016). Each exogenous variable $\widehat{U} \in \widehat{U}$ is a 2-dimensional vector, with each dimension sampled independently from a uniform distribution between -1 and 1. The discriminator is a feedforward neural network with 3 hidden layers of width 64. The GAN-NCM for the third approach with abstractions has the exact same parameter settings but is modeled over \mathbf{V}_H instead. Consequently, the NCM for the third approach has fewer parameters since it only requires three functions for \mathbf{V}_H , compared to the six functions for \mathbf{V}_L .

The GAN-NCMs are trained in the style of Wasserstein GANs (Arjovsky, Chintala, and Bottou 2017), where the objective is to minimize the Earth-Mover distance via the Kantorovich-Rubenstein duality (Villani 2009):

$$\min_G \max_{D \in \mathcal{D}_D} \mathbb{E}_{\mathbf{x} \sim P_r} [D(\mathbf{x})] - \mathbb{E}_{\tilde{\mathbf{x}} \sim P_g} [D(\tilde{\mathbf{x}})], \quad (98)$$

where G is a generating model (e.g. the NCM \widehat{M}), D is a discriminatory model, also called a critic (not to be confused with the variable $D \in \mathbf{V}_L$), \mathcal{D}_D is the set of 1-Lipshitz functions, P_r is a real distribution (i.e. from \mathcal{M}_L), and P_g is the distribution induced by G .

¹⁸In fact, this second approach can also be considered abstracting the space of \mathbf{V}_L . Specifically, the intervariable and intravariabile clusters are all singleton clusters, but the values of \mathbf{V}_H are renamed such that optimization is easier (similar to the ideas discussed in Sec. 4).

For identification experiments, models were trained for 1000 epochs on datasets with $n = 10^4$ samples. 10 trials were performed with each approach, with 4 reruns for each trial for hypothesis testing purposes. In a single run, two parameterizations of the NCM are initialized with one aiming to minimize the query and one aiming to maximize it. In each iteration, a batch of real data is provided, and a batch of fake data is generated by the NCM. Given these two batches, the discriminator is trained to minimize the loss following Eq. 98:

$$L_D = \mathbb{E}_{\tilde{\mathbf{x}} \sim P_g} [D(\tilde{\mathbf{x}})] - \mathbb{E}_{\mathbf{x} \sim P_r} [D(\mathbf{x})],$$

where P_g and P_r refer to the fake and real datasets respectively. In words, the loss is computed by taking the expected score of the critic on fake samples subtracted by the expected score of the critic on real samples, indicating better performance if the critic gives higher scores to real samples. After each training iteration, the gradients of the discriminator D are clamped between $[-0.01, 0.01]$ to enforce the Lipschitz constraint.

Following an iteration of the discriminator, another batch of fake data is sampled from the generator (NCM \widehat{M}), and the weights of the generator are updated with the loss

$$L_G = -\mathbb{E}_{\tilde{\mathbf{x}} \sim P_g} [D(\tilde{\mathbf{x}})] + \lambda L_Q. \quad (99)$$

The first term is the expected score of the critic on the fake samples, which should be maximized by the generator to create as convincing samples as possible. The second term is a query loss, intended to push the model to simultaneously maximize or minimize the query. In practice, this is done by calculating the distance between the intended value of the query and query samples from the generator using some distance function. For example, for the GAN-NCM in the first non-abstraction approach, L_Q is defined as:

$$L_Q(\hat{\mathbf{b}}) = \pm \left(\frac{1}{|\hat{\mathbf{b}}|} \sum_{\hat{b} \in \hat{\mathbf{b}}} \hat{b} - 25 \right),$$

where $\hat{\mathbf{b}}$ is a batch of samples from $P^{\widehat{M}}(B_{D=d})$ computed from Xia, Pan, and Bareinboim (2023, Alg. 2). If this quantity, which is simply a mean over the batch samples, is maximized (resp. minimized), then that would also maximize (resp. minimize) the query $Q = P^{\widehat{M}}(B_{D=d} \geq 25)$. For the second approach with normalized data, the 25 is not subtracted as it is already centered around 0. For the third approach working in the abstracted space, the log loss is calculated instead, since values of B_H are binary. λ is a hyperparameter to indicate the strength of the query loss term; in our experiments it was set to 10^{-4} and decreased logarithmically to 10^{-8} by the end of training.

For the visualization of the results in Fig. 10a, the query is estimated from both the model which optimized to maximize it (denote as Q_{\max}) and the model which optimized to minimize it (Q_{\min}). Since the query is identifiable (see Sec. C.5), we expect $Q_{\max} - Q_{\min} = 0$ under perfect optimization. However, as optimization is not perfect, a hypothesis testing procedure must be used to check if $Q_{\max} - Q_{\min} < \varepsilon$ for

some threshold ε . As suggested by Xia, Pan, and Bareinboim (2023), we rerun each trial 4 times and take the upper 95% confidence bound of the mean of $Q_{\max} - Q_{\min}$ from the 4 reruns. Then the means of these upper confidence bounds across 10 trials are plotted in Fig. 10a with 95% confidence intervals.

For estimation experiments, models were trained for 200 epochs on datasets. 10 trials were performed for each approach and each setting of sample size varying logarithmically from $n = 10^3$ to 10^5 samples. The training procedure is performed identically to the identification experiments, except only one parameterization is trained, and the query loss in Eq. 99 is not added. After training, queries are estimated from each of the models using Xia, Pan, and Bareinboim (2023, Eq. 4) with 10^5 Monte Carlo samples, and they are compared with the ground truth value calculated from \mathcal{M}_L as described in App. C.1. The mean absolute error (MAE) is computed between the two values and plotted in Fig. 10b with 95% confidence intervals across the 10 trials for each of the sample size settings.

All NCMs are trained with a learning rate of 10^{-4} , and discriminators are trained with a learning rate of 2×10^{-4} . Models are optimized with the RMSProp optimizer (Hinton, Srivastava, and Swersky 2018), recommended by the WGAN paper. Estimation experiments are performed with a batch size of 128, and identification experiments are performed with a batch size of 1000 (larger size for more representative sample is important in identification). All feedforward networks are initialized with Glorot initialization (Glorot and Bengio 2010). Hyperparameter tuning was done by hand with the suggestions from referenced sources. Similar hyperparameters did not have noticeable effects on performance, so more rigorous hyperparameter tuning was not conducted.

C.3 Colored MNIST Experimental Setup

The Colored MNIST experiment in Sec. 5.2 is performed on a modified version of the MNIST dataset of handwritten digits (Deng 2012). The setting is modeled over variables \mathbf{V}_L consisting of a digit label D , a color label C , and the pixels consisting of a $3 \times 32 \times 32$ MNIST image with color channels. Both D and C take integer values from 0 to 9, formatted as one-hot vectors in the data. However, the mechanisms for which the image is generated is unknown, since we do not know all of the details of how humans handwrite digits. Instead, we directly work on the high level space of variables $\mathbf{V}_H = \{D, C, I\}$, obtained by clustering all of the pixels into one variable, called image I . Samples from the observational distribution $P(\mathbf{V}_L)$ are generated using the following approach:

1. A sample is drawn from exogenous variable U_{CD} , which takes values from 0-9, indicating what the intended digit is.
2. With probability 0.85, set $C \leftarrow U_{CD}$. Otherwise, choose uniformly at random from the 10 values. Similarly, but independently from C , set $D \leftarrow U_{CD}$ with probability 0.85, otherwise choose uniformly at random.
3. Given C and D , sample an image from the MNIST dataset

with label D , then color the digit with the color corresponding to C on the gradient in Fig. 11b.

The causal diagram \mathcal{G} over \mathbf{V}_L is unknown because it is unclear how individual pixels are related. However, the C-DAG \mathcal{G}_C over \mathbf{V}_H is shown in Fig. 11a, which is compatible with the data generating process mentioned above. Specifically, image I is caused by color C and digit D , which are highly correlated through unobserved confounding.

The task is to train a model \widehat{M}_H over variables \mathbf{V}_H constrained by the graph \mathcal{G}_C such that \widehat{M}_H induces the distribution $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$ (i.e. it is perfectly trained to match the observational data sampled from $P^{\mathcal{M}_L}(\mathbf{V}_L)$), and then use it to produce realistic digit samples from three different causal queries:

1. $P(I \mid D = 0)$: the distribution of images conditional on digit = 0. In the dataset, the digit 0 is highly correlated with the color red, so samples from this distribution should be images of handwritten 0s, most of which are red.
2. $P(I_{D=0})$: the distribution of images when intervened on digit = 0. When an intervention is performed, spurious correlations are ignored. The color is sampled like normal, but then the digit is forced to become 0 regardless of the color. Hence, samples from this distribution should be images of handwritten 0s but with colors evenly distributed.
3. $P(I_{D=0} \mid D = 5)$: the counterfactual distribution of images of what they would have been had digit been forced to be 0 given that the digit was originally 5. When conditioning on $D = 5$, the samples are filtered such that only ones with $D = 5$ remain, but then these samples are intervened and forced to take the digit 0 instead. Consequently, samples from this distribution should be images of handwritten 0s that retain the color of the 5s, which are typically cyan.

Three different approaches are compared:

1. The first is a basic conditional GAN that learns the correlation between digit D and image I . The conditional GAN ignores the information in \mathcal{G}_C and therefore is naïve to the causal invariances in the data.
2. The second is a GAN-NCM (Xia, Pan, and Bareinboim 2023) that is constrained by \mathcal{G}_C and is directly fitted on the data $\tau(P^{\mathcal{M}_L}(\mathbf{V}_L))$. In this case τ simply clusters the pixels together into the image I , but the mapping between domains of \mathbf{V}_L and \mathbf{V}_H is the identity mapping. In other words, the intravariation clusters \mathbb{D} can be thought of as the singleton partition of all domains, and as a result the space of \mathbf{V}_H is identical to \mathbf{V}_L .
3. The third is a GAN version of the RNCM from Sec. 4, called GAN-RNCM. The GAN-RNCM is also constrained by \mathcal{G}_C but learns its own abstraction function $\widehat{\tau}$. Specifically, in a typical instantiation of the \mathcal{G}_C -RNCM in this case, we would have $\widehat{\tau} = (\widehat{\tau}_C, \widehat{\tau}_D, \widehat{\tau}_I)$, where each subfunction learns a mapping to a representation space (akin to learning intravariation clusters). For this experiment, we only parameterized $\widehat{\tau}_I$, since C and D are low-dimensional and are already easy to learn. $\widehat{\tau}_I$ is trained to map to a space that preserves bijectivity (as demanded by

Prop. 5) as well as maximizing information retained about C and D . See the next subsection for specific details.

The results are illustrated in Fig. 5.2. The GAN-RNCM clearly outperforms the other two approaches, and we even observed a shorter runtime. Although the GAN-NCM is, in theory, supposed to be able to capture the intended distributions, we believe its failure is a result of the difficulty of simultaneously optimizing two different tasks: (1) image generation is already a challenging task with a lot of attention in the deep learning community, and (2) learning a joint distribution with causal constraints is also challenging. The GAN-RNCM breaks the problem into two parts. The representation learning of $\hat{\tau}$ solves the problem of dealing with high dimensional images, reducing the space to a much simpler space in which the distribution with causal constraints can be learned more easily.

C.4 Colored MNIST Models and Hyperparameters

The GAN-NCM used in the Colored MNIST experiment is different from the ones used in the Nutrition experiment since it is fitted on a different graph, specifically \mathcal{G}_C from Fig. 11a. Since the function for the image, \hat{f}_I , is responsible for generating an entire image, we leverage the technology of convolutional neural networks to produce higher quality results. Specifically, we use the state-of-the-art research on conditional image generation implemented by Brock, Donahue, and Simonyan (2019), called BigGAN. \hat{f}_I is designed by first mapping the inputs, color C and digit D , through a feedforward neural network to an internal representation, which is then piped into the 32×32 image-size architecture with 64 feature maps implemented by the BigGAN authors. The functions \hat{f}_C and \hat{f}_D are simply feedforward neural networks. In this model, all feedforward nets have 3 hidden layers, with widths that depend on the size of the inputs and outputs using the formula $2 \times i \times o$, where i is the total dimensionality of all endogenous and exogenous inputs, and o is the output dimensionality (number of channels for images). Each exogenous variable $\hat{U} \in \hat{\mathbf{U}}$ is a δ -dimensional vector, where δ is the sum of the dimensions of all variables in the confounded clique represented by \hat{U} , and each dimension is sampled independently from a uniform distribution between -1 and 1. The discriminator first pipes image inputs through a deconvolutional component like implemented in BigGAN, before combining the internal representation with other variables to pipe through a feedforward neural network with 3 hidden layers of width 128. Layer normalization is applied between layers of feedforward nets, and batch normalization (Ioffe and Szegedy 2015) is applied between convolutional layers.

The conditional GAN approach is implemented similarly, but without \hat{f}_C . Training is done identically to the Wasserstein GAN approach in the Nutrition experiment (described in App. C.2), but without the query loss in Eq. 99, as identification is not performed.

The GAN-RNCM $\langle \hat{\tau}, \hat{M} \rangle$ is trained in a two part procedure, first training $\hat{\tau}$ and then training \hat{M} on the space

defined by $\hat{\tau}$. Only the abstraction function for the image, $\hat{\tau}_I$, is trained for this experiment. It is modeled in two parts: (1) a convolutional neural network with three convolutional layers (with 64, 128, and 256 feature maps respectively) mapping the image to a 128-dimensional vector, and (2) a feedforward neural network with 3 hidden layers of width 128 mapping the convolutional output to a 64-dimensional representation space.

$\hat{\tau}_I$ is trained for 500 epochs. In each epoch, a batch of the colored MNIST digits is sampled and passed through $\hat{\tau}$ to obtain a representation. Then, $\hat{\tau}$ is trained with the loss

$$L_{\tau_G}(\hat{i}_H, i_L, c_L, d_L) = \left\| \tau^{-1}(\hat{i}_H; \theta_{\tau^{-1}}) - i_L \right\|^2 \quad (100)$$

$$+ \lambda_g L_g(g(\hat{i}_H; \theta_g), c_L, d_L), \quad (101)$$

where \hat{i}_H is the 64-dimensional representation output from $\hat{\tau}_H(i_L)$; and i_L, c_L, d_L are the original data points¹⁹. The first term (Eq. 100) is a reconstruction loss that ensures that Prop. 5 holds. τ^{-1} is another neural network (parameterized by $\theta_{\tau^{-1}}$) in the style of BigGAN that upscales the representation \hat{i}_H back to an image of size $3 \times 32 \times 32$. The term is simply the MSE of the reconstruction with the original image, ensuring that both the encoder τ and decoder τ^{-1} are trained to be able to reproduce the input. Later, when sampling images of I , τ^{-1} is also used to reconstruct image samples. The second term (Eq. 101) is a classification loss added to improve the learned representation to differentiate between different values of c and d . g is feedforward neural network parameterized by θ_g with 3 hidden layers of width 128, which outputs a prediction for c and d given the representation \hat{i}_H . Any classification loss can be used for L_g , and we choose binary cross-entropy loss since C and D are one-hot vectors. λ_g is a regularization term which takes a value of 0.1 in our experiments.

After training $\hat{\tau}_I$, the NCM \hat{M} of the RNCM is trained similarly to the other approaches, but it is instead trained on top of $\hat{\tau}(\mathbf{V}_L)$ instead of \mathbf{V}_L . That is, instead of outputting an image from \hat{f}_I , it outputs a 64-dimensional real vector, representing $\hat{\tau}_I(I)$. Hence, in the RNCM, \hat{f}_I is replaced with a feedforward neural network with 3 hidden layers of width $2 \times i \times o$, as with the other functions.

For the training of $\hat{\tau}_I$, as well as the training of all three generative models, 10^5 samples are provided in the dataset $P(\mathbf{V}_L)$. All models are trained with a learning rate of 10^{-4} , and discriminators are trained with a learning rate of 2×10^{-4} . GAN models are optimized with the RMSProp optimizer, and the training procedure for $\hat{\tau}_I$ is optimized with Adam (Kingma and Ba 2015). All training is performed with a batch size of 128. All feedforward networks are initialized with Glorot initialization. Hyperparameter tuning was done by hand with the suggestions from referenced sources. Similar hyperparameters did not have noticeable effects on performance, so more rigorous hyperparameter tuning was not conducted.

Samples from the three competing approaches, as well as the original data generating ground truth, are collected

¹⁹Although C and D are not part of the cluster with I , they can still be used in the training process for $\hat{\tau}_I$ as long as they are not used as inputs to $\hat{\tau}_I$.

from the three queries discussed in Sec. C.3 and compared in Fig. 12. Queries are sampled via Xia, Pan, and Bareinboim (2023, Alg. 1).

C.5 Proofs of Identifiability

In this section, we show proofs that the queries in the experiments are identifiable, leveraging do-calculus (denote R1, R2, and R3 as the three rules) and counterfactual axioms (Pearl 2000).

Proposition 6. $P(B_{D=d} = b)$ is identifiable from $P(\mathbf{V}_L)$ and \mathcal{G} from Fig. 5 (left). ■

Proof.

$$\begin{aligned} P(B_{D=d} = b) &= \sum_r P(b_d | r_d)P(r_d) \\ &= \sum_r P(b | r, d)P(r_d) && \text{R2} \\ &= \sum_r P(b | r, d)P(r) && \text{R3,} \end{aligned}$$

and the final result can be computed from observational $P(\mathbf{V}_H)$ as there are no more interventional terms. □

Proposition 7. $P(B_{H[D_H=d]} = b)$ is identifiable from $P(\mathbf{V}_H)$ and \mathcal{G}_C from Fig. 5 (right). ■

Proof.

$$\begin{aligned} P(B_{H[D_H=d]} = b) &= \sum_z P(z_d)P(b_d | z_d) \\ &= \sum_z P(z | d)P(b_d | z_d) && \text{R2} \\ &= \sum_z P(z | d)P(b_{dz}) && \text{R2} \\ &= \sum_z P(z | d)P(b_z) && \text{R3} \\ &= \sum_z P(z | d) \sum_{d'} P(b_z | d'_z)P(d'_z) \\ &= \sum_z P(z | d) \sum_{d'} P(b | d', z)P(d'_z) && \text{R2} \\ &= \sum_z P(z | d) \sum_{d'} P(b | d', z)P(d') && \text{R3,} \end{aligned}$$

and the final result can be computed from observational $P(\mathbf{V}_H)$ as there are no more interventional terms. □

Proposition 8. $P(I = i | D = d)$ is identifiable from $P(\mathbf{V}_H)$ and \mathcal{G}_C from Fig. 11a. ■

Proof. This result is trivial, as $P(I = i | D = d)$ is an observational quantity and can therefore be computed as

$$\begin{aligned} P(I = i | D = d) &= \frac{P(i, d)}{P(d)} = \frac{\sum_{c'} P(i, c', d)}{\sum_{i', c'} P(i', c', d)} = \frac{\sum_c P(\mathbf{v}_H)}{\sum_{i, c} P(\mathbf{v}_H)}. \end{aligned}$$

□

Proposition 9. $P(I_{D=d} = i)$ is identifiable from $P(\mathbf{V}_H)$ and \mathcal{G}_C from Fig. 11a. ■

Proof.

$$\begin{aligned} P(I_{D=d} = i) &= \sum_c P(i_d | c_d)P(c_d) \\ &= \sum_c P(i | c, d)P(c_d) && \text{R2} \\ &= \sum_c P(i | c, d)P(c) && \text{R3,} \end{aligned}$$

and the final result can be computed from observational $P(\mathbf{V}_H)$ as there are no more interventional terms. □

Proposition 10. $P(I_{D=d} | D = d')$ is identifiable from $P(\mathbf{V}_H)$ and \mathcal{G}_C from Fig. 11a. ■

Proof.

$$\begin{aligned} P(I_{D=d} = i | D = d') &= \frac{P(i_d, d')}{P(d')} \\ &= \frac{\sum_c P(i_d, d', c)}{P(d')} \\ &= \frac{\sum_c P(i_{dc}, d', c)}{P(d')} && \text{C1} \\ &= \frac{\sum_c P(i_{dc})P(d', c)}{P(d')} && \text{C2} \\ &= \frac{\sum_c P(i | d, c)P(d', c)}{P(d')} && \text{R2} \\ &= \sum_c P(i | d, c)P(c | d') \end{aligned}$$

where ‘‘C1’’ refers to the counterfactual axiom of composition, and ‘‘C2’’ refers to the C-factor decomposition of counterfactual variables (Correa, Lee, and Bareinboim 2021). The final result can be computed from observational $P(\mathbf{V}_H)$ as there are no more interventional terms. □

C.6 Hardware

All models were trained on NVIDIA Tesla V100 GPUs provided by Amazon Web Services, totalling approximately 2000 GPU hours for the final results.

D Further Discussion

D.1 The Constitution Hierarchy and Learning Intervariable Clusters

This section provides a detailed discussion on intervariable clusters \mathbb{C} , on the relationship of clusters at different levels of granularity, and best practices on how to choose the right level of granularity when several options are available.

On the intervariable level, one could cluster several lower level variables together and call the cluster a variable itself. The low-level variables do not *cause* the high level variable,

but rather they *constitute* it (as discussed in Chalupka, Perona, and Eberhardt (2015), and similar to discussions in probabilistic relational models on part-whole relationships (Winston, Chaffin, and Herrmann 1987)). In other words, they are two interpretations of the same content. The difference is important: the relationship is bidirectional, and one cannot directly intervene on one without simultaneously intervening on the other.

When comparing variables at different levels of granularity, a hierarchy arises, which we call the *constitutional hierarchy*. That is, any phenomenon can be viewed across another dimension that determines the level of granularity of interpreting the variables. This is illustrated in Fig. 15. Variables can be organized by what constitutes what, and, in a proper abstraction, causal properties should be preserved across different levels of granularity. Generally speaking, most studies focus on one specific level of abstraction, so the task of choosing the most appropriate level of granularity can be important. This is precisely the problem of choosing intervariable clusters, as the coarseness of the clusters induces a natural interpretation of the variables.

If the data scientist finds themselves in a situation where they have to construct the causal diagram themselves, or if a provided causal diagram is not at the right level of abstraction, then refining the abstraction level becomes a nontrivial task. The decision of whether to cluster variables together depend on various factors, which we list and elaborate in the sequel:

- C1. *Non-causal* relationships should not be visible.
- C2. The resulting clustering should be admissible.
- C3. The queries of interest should be answerable.
- C4. The queries of interest should be identifiable.
- C5. The result should be as coarse as possible.

Each of these conditions can be formalized to enable a systematic discussion on how to choose an appropriate intervariable clustering.

Condition C1. When considering models at extremely low levels of abstraction, there may be too much detail to properly label every relationship as a causal one.²⁰ For example, like in the top level of Fig. 15, at the atom or molecule level, there are interactions between particles that are studied in the physical sciences such as bonds. That is, two particles may have linked behavior, but it is not accurate to call such a relationship causal. As another example, suppose we consider images at the pixel level. In many tasks, one may be interested in the local dependences between pixels, and it may therefore be reasonable to model pixels using an undirected model (e.g., like a Markov random field). Once again, the relationship between the pixels may not necessarily be considered causal.

We note importantly that “non-causal” in this context does not refer to spurious effects, nor does it mean anti-causal in

²⁰We do not provide an exact definition of a causal relationship, as this is a deeply philosophical topic that is out of the scope of this work. Still, we acknowledge that there may be cases where two low-level variables are related in a way that is not well-defined with respect to interventions in an SCM. We use the term “non-causal” as a bucket term for all such cases and make no assumptions about the natures of these relationships.




Figure 15: A visual example of the constitution hierarchy. Green edges indicate causal relations, blue edges indicate constitutional relations, and black edges are undefined, low-level relationships. The same objects are shown at three levels of granularity: molecules (top), parts (middle), and complete objects (bottom). A causal relationship might indicate that a car is the cause of Bob’s injury, but it would be inaccurate to say that the molecules of a car “cause” the car. A particular study should focus on the variables from one specific level of granularity, and choosing the most appropriate one is the goal of the problem of choosing intervariable clusters.

the sense that the direction of causality is reversed, as these are still well-defined from a causal perspective. For example, in a scenario modeled by Fig. 16, the relationship between X and Y is not considered a non-causal relationship. Despite Y not causing X , their relationship is still considered causal due to the cause from X from Y . Further, although there is spurious correlation through the unobserved confounding, this is still considered a causal type of relationship since there is some causal effect from the unobserved confounder to X and Y . “Non-causal” relationships refer to relationships between variables that cannot be defined in a causal manner. When considering non-causal relationships, the SCM framework is no longer compatible since it is not defined how to model such relationships. Nonetheless, the framework presented in this paper can still be used as long as C1 is satisfied.

The term “non-causal” is very broad and could encompass many different types of relationships. Studying specific types of non-causal relationships is out of the scope of this work, so we will treat these relationships in the same way. For every subset of variables (say $\mathbf{Z} \subseteq \mathbf{V}_L$) related through some non-causal relationship, we will assume that there is some function $f_{\mathbf{Z}}$ that takes as input variables from \mathbf{U}_L and \mathbf{V}_L and maps it jointly to the space of \mathbf{Z} . We note that this is a general encoding of such relationships that makes minimal assumptions. For example, perhaps the relationship between three variables, X , Z , and Y can be described using a Markov random field like shown in Fig. 17, where $X \perp\!\!\!\perp Y \mid Z$. Our assumption states that the behavior of these three variables can be abstracted into one function $f_{\{X,Z,Y\}}$ which outputs values of X , Z , and Y , therefore losing the independence

information. (It is possible, however, that these could be implicitly encoded through the intravariabile clusters). These non-causal relationships can be expressed graphically.

Definition 21 (Non-causal Graph). Let \mathbf{V}_L be a set of variables and $\bar{\mathcal{G}} = \langle \mathbf{V}_L, \mathbf{E} \rangle$ be an undirected graph with nodes representing \mathbf{V}_L . Then, an edge (V_1, V_2) is in \mathbf{E} if and only if there is some non-causal relationship between V_1 and V_2 . ■

We will assume that, in settings which intervariable clusters must be learned, we are given knowledge of the existence of non-causal relationships through a non-causal graph of the lowest level variables \mathbf{V}_L . This information could be acquired simply as an assumption of the user, as they are typically only present at the lowest levels of abstraction (e.g. the relationship between pixels in an image or between atoms of an object). Under this assumption, distributions over \mathbf{V}_L can be treated as if variables in the same connected component in $\bar{\mathcal{G}}$ share a common cause, and they can be factorized as such.

Condition C2. By definition of admissibility, two variables should not be in the same intervariable cluster if doing so forms a cycle in the order of the functions or graph. This condition can be verified given causal diagram \mathcal{G} of the low level variables \mathbf{V}_L .

Condition C3. C3 and C4 depend on the user’s needs based on the queries of interest, \mathcal{Q} , in downstream tasks. The ideal level of abstraction can be determined by the groupings of variables in these queries. C3 essentially enforces that the level of abstraction should be kept low enough to retain the ability to differentiate the nuances between individual variables of interest in the queries. For example, if one would like to study the causal effect of a drug X on recovery rate Y , but Y was clustered with another variable Z representing blood pressure, then it no longer becomes possible to answer queries specifically about Y without including Z . This is formally defined below.

Definition 22 (Cluster Answerability). A counterfactual query $Q = P^{\mathcal{M}_L}(\mathbf{Y}_{1[x_1]}, \mathbf{Y}_{2[x_2]}, \dots)$ is answerable from intervariable clusters \mathcal{C} if and only if for all \mathbf{Y}_i (and \mathbf{X}_i), there exists $\mathcal{C}_i \subseteq \mathcal{C}$ such that $\mathbf{Y}_i = \bigcup_{\mathcal{C} \in \mathcal{C}_i} \mathcal{C}$. ■

In other words, all queries should be written in terms of unions of clusters.

Condition C4. Even if answerability is not violated, clustering variables together or projecting them out will always result in a loss of information. Such a loss may affect the identifiability of the queries. Ideally, one would not want to drop information that would change the status of an already identifiable query, and this is enforced by C4.

Condition C5. Finally, C5 is usually desirable following the idea of Occam’s razor, i.e., all else being equal, simplicity should be preferred. We define coarseness as follows.

Definition 23 (Coarseness). Let \mathcal{C}_1 and \mathcal{C}_2 be two intervariable clusterings of \mathbf{V}_L . \mathcal{C}_1 is said to be coarser than \mathcal{C}_2 (equivalently, \mathcal{C}_2 is finer than \mathcal{C}_1) if and only if for every $\mathcal{C}_2 \in \mathcal{C}_2$, either $\mathcal{C}_2 \cap \bigcup_{\mathcal{C}_1 \in \mathcal{C}_1} \mathcal{C}_1 = \emptyset$, or there exists $\mathcal{C}_1 \in \mathcal{C}_1$ such that $\mathcal{C}_2 \subseteq \mathcal{C}_1$.




Figure 16: Example of spurious relationships.




Figure 17: Example of non-causal graph.

With this definition, the goal is to find a maximally coarse clustering \mathcal{C} that satisfies conditions 1-4. That is, there should not exist a coarser clustering \mathcal{C}' that also satisfies the conditions, although the coarsest clustering is not necessarily unique.

It turns out that the set of all possible maximally coarse clusterings can be “bounded” in a sense using C1, C2, and C3. The “minimally” coarse clustering can be described with the following lemma.

Lemma 7. Let $\bar{\mathcal{G}}$ be a non-causal graph over \mathbf{V}_L . Let \mathcal{C}_{\min} be the intervariable clustering of \mathbf{V}_L composed of the connected components of $\bar{\mathcal{G}}$. Then, any clustering \mathcal{C}' violates condition C1 if and only if it is not equal to or coarser than \mathcal{C}_{\min} . ■

Proof. We first note that \mathcal{C}_{\min} does not violate condition C1 because by construction, if there exists a noncausal edge between V_1 and V_2 in $\bar{\mathcal{G}}$, then they must be in the same cluster. Any coarser clustering has the same property.

If \mathcal{C}' is not equal to or coarser than \mathcal{C}_{\min} , that means that there exists $\mathcal{C} \in \mathcal{C}_{\min}$ such that $\mathcal{C} \not\subseteq \mathcal{C}'$ for all $\mathcal{C}' \in \mathcal{C}'$. This implies that there exists $V_1, V_2 \in \mathcal{C}$ such that belong to different clusters \mathcal{C}'_1 and \mathcal{C}'_2 in \mathcal{C}' . However, since V_1 and V_2 were in the same cluster in \mathcal{C}_{\min} , there must exist some non-causal path from V_1 to V_2 in $\bar{\mathcal{G}}$. This means that there is a noncausal connection between clusters \mathcal{C}'_1 and \mathcal{C}'_2 , so \mathcal{C}' violates condition C1. □

Intuitively, the minimally coarse clustering must at least cover the connected components of $\bar{\mathcal{G}}$ to abstract away the non-causal relations and satisfy condition C1. Additionally, admissibility (C2) adds another constraint.

Lemma 8. Let \mathcal{C} be a set of intervariable clusters that are not admissible w.r.t. \mathcal{G} . Let \mathcal{C}^* be the set of clusters that is coarser than \mathcal{C} such that \mathcal{C}_1 and \mathcal{C}_2 in \mathcal{C} are merged in \mathcal{C}^* if and only if \mathcal{C}_1 and \mathcal{C}_2 are in a cycle in $\mathcal{G}_{\mathcal{C}}$. Then, any set of clusters that are coarser than \mathcal{C} and are admissible w.r.t. \mathcal{G} are equal to or coarser than \mathcal{C}^* . ■

Proof. First, we note that \mathcal{C}^* is admissible w.r.t. \mathcal{G} by construction, since all cycles have been merged. If \mathcal{C}' is a clustering that is coarser than \mathcal{C} but not \mathcal{C}^* , then there must exist $\mathcal{C}^* \in \mathcal{C}^*$ such that $\mathcal{C}^* \not\subseteq \mathcal{C}'$ for all $\mathcal{C}' \in \mathcal{C}'$. Then, there must exist $V_1, V_2 \in \mathcal{C}^*$ such that V_1 and V_2 are in different clusters in $\mathcal{C}_1, \mathcal{C}_2 \in \mathcal{C}'$. However, this implies that there is still a cycle between \mathcal{C}_1 and \mathcal{C}_2 in $\mathcal{G}_{\mathcal{C}'}$, breaking admissibility. □

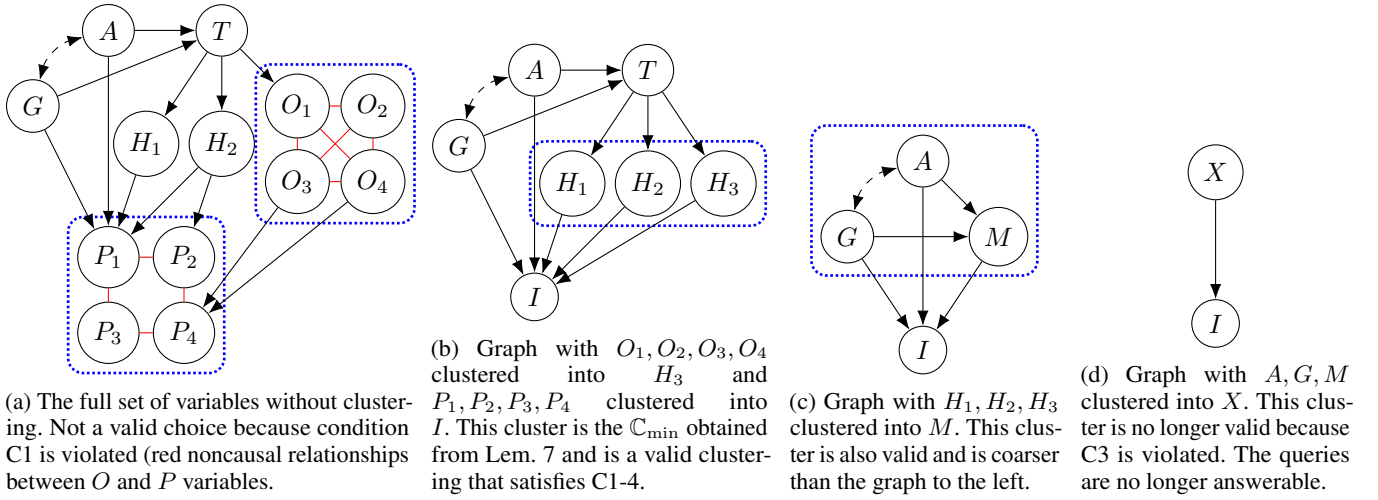


Figure 18: Graphs for Example 16. Blue outline illustrates changes in clustering between each graph.

On the other hand, the maximally coarse clustering is related to the concept of maximally answerable clusters, defined below.

Definition 24 (Maximally Answerable Clusters). Let \mathcal{Q} be a set of counterfactual queries over variables \mathbf{V}_L , and let $\mathbb{V}_{\mathcal{Q}}$ be the set of all subsets of \mathbf{V}_L that are used in a term in \mathcal{Q} (i.e. for every counterfactual term $\mathbf{Y}_{\mathbf{x}}$ in any of the queries, both \mathbf{Y} and \mathbf{X} should be included in $\mathbb{V}_{\mathcal{Q}}$). Denote $\mathbb{V}_{\mathcal{Q}}^*$ as the extension of $\mathbb{V}_{\mathcal{Q}}$ closed under intersection, that is, $\mathbb{V}_{\mathcal{Q}} \subseteq \mathbb{V}_{\mathcal{Q}}^*$, and $V_1, V_2 \in \mathbb{V}_{\mathcal{Q}}^* \Rightarrow V_1 \cap V_2 \in \mathbb{V}_{\mathcal{Q}}^*$. Then, an intervariable clustering \mathcal{C}_{\max} is called the maximally answerable cluster w.r.t. \mathcal{Q} if and only if \mathcal{C}_{\max} is the set of atoms of $\mathbb{V}_{\mathcal{Q}}^*$, or in other words, $\mathbf{V}_i \in \mathcal{C}_{\max}$ if and only if $\mathbf{V}_i \in \mathbb{V}_{\mathcal{Q}}^*$ and there exists no $\mathbf{V}_j \in \mathbb{V}_{\mathcal{Q}}^*$ such that $\mathbf{V}_j \subsetneq \mathbf{V}_i$. ■

The concept of maximally answerable clusters can be best shown through an example.

Example 15. Suppose

$$\mathbf{V}_L = \{A, B, C, D, E, F, G, H, I, X, Z\}$$

and the queries of interest $\mathcal{Q} = \{Q_1, Q_2, Q_3\}$ are as follows:

$$\begin{aligned} Q_1 &= P(\{a, b, c, d, e\}_x) \\ Q_2 &= P(\{c, d, e, f, g, h\}_{x,z}) \\ Q_3 &= P(\{e, f, h, i, z\}_x). \end{aligned}$$

We first dissect the terms in the queries, including the subscripts, to obtain $\mathbb{V}_{\mathcal{Q}}$. Doing so yields

$$\mathbb{V}_{\mathcal{Q}} = \{\{A, B, C, D, E\}, \{C, D, E, F, G, H\}, \{E, F, H, I, Z\}, \{X\}, \{X, Z\}\}.$$

These sets are shown at the top of Fig. 19. Note that many of the sets overlap. The set of nonoverlapping subregions becomes the maximally answerable clusters, as shown at the bottom of Fig. 19. In this case, this means

$$\mathcal{C}_{\max} = \{\{A, B\}, \{C, D\}, \{E\}, \{F, H\}, \{G\}, \{I\}, \{X\}, \{Z\}\}.$$

Note that no set in \mathcal{C}_{\max} overlaps with each other, but all sets are subsets of some set in $\mathbb{V}_{\mathcal{Q}}$. ■

It turns out that the set of maximally answerable clusters is also the maximally coarse clustering that does not violate condition C3, as shown below.

Lemma 9. Let \mathcal{C}_{\max} be the intervariable clustering of \mathbf{V}_L that is the maximally answerable cluster with respect to a set of counterfactual queries \mathcal{Q} . Then, any clustering \mathcal{C}' violates condition C3 if and only if it is not equal to or finer than \mathcal{C}_{\max} . ■

Proof. We first note that \mathcal{C}_{\max} does not violate condition C3, since if there existed a query with a term set \mathbf{Z} such that \mathbf{Z} is not a union of clusters of \mathcal{C}_{\max} , then either there exists $Z \in \mathbf{Z}$ such that Z is not in any set of \mathcal{C}_{\max} , or $Z \in \mathcal{C}$ for some $\mathcal{C} \in \mathcal{C}_{\max}$ such that $\mathcal{C} \subsetneq \mathbf{Z}$. However, the first case is impossible since if $Z \in \mathbf{Z}$, then Z must be added to some set in $\mathbb{V}_{\mathcal{Q}}$. The second case is not possible either because then $\mathbf{Z} \cap \mathcal{C} \neq \emptyset$, which means that the clusters of \mathcal{C}_{\max} could be divided further, contradicting its definition. If \mathcal{C}_{\max} does not violate condition C3, then no finer clustering can violate it either because a union of clusters in \mathcal{C}_{\max} can always translate to a union of clusters in a finer clustering.

If \mathcal{C}' is not equal to or finer than \mathcal{C}_{\max} , then there exists at least one cluster $\mathcal{C}' \in \mathcal{C}'$ such that $\mathcal{C}' \subsetneq \mathcal{C}$ for all $\mathcal{C} \in \mathcal{C}_{\max}$. This means that there exist $V_1, V_2 \in \mathcal{C}'$ such that V_1 and V_2 are not in the same cluster in \mathcal{C}' . This implies that there exists at least one query such that one of its terms contains one of V_1 or V_2 but not the other, in which case, such a term cannot be described as a union of clusters from \mathcal{C}' , violating condition C3. □

Intuitively, we cannot simply cluster all variables together because we may want to answer queries that require separate consideration of different sets of variables. At best, we can only cluster together terms that always appear together in all queries. Additionally, we note that finding the maximally answerable clusters can be done in polynomial time. One simple approach is to add one set of $\mathbb{V}_{\mathcal{Q}}$ at a time, repeatedly computing the intersection with the existing clusters. At any point, there can be at most $n = |\mathbf{V}_L|$ clusters.

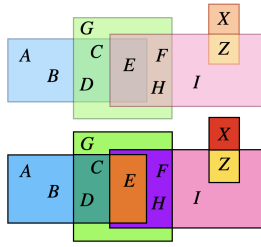


Figure 19: **Top:** A venn diagram of all of the variables and what set of \mathbb{V}_Q to which they belong. **Bottom:** The collection of nonoverlapping intersections of all sets that form the maximum answerable clusters.

With these points in mind, we provide Alg. 3 for learning intervariable clusters. We note that Alg. 3 returns a maximally course clustering satisfying all conditions.

Proposition 11. *Alg. 3 returns a maximally course clustering that satisfies C1-4, or returns FAIL if one does not exist. ■*

Proof. We first note that if FAIL is returned, then there does not exist a clustering that satisfies all four conditions. Notably, if \mathbb{C}_{\max} is not coarser than \mathbb{C}_{\min} at the beginning, then FAIL is returned, as a consequence of Lemmas 7, 8, and 9. FAIL is also returned if \mathbb{C}_{\min} does not satisfy identifiability of all queries, since further clustering will only remove edges and nodes and cannot make a non-ID query turn ID.

If \mathbb{C} is returned at the end, then \mathbb{C} must satisfy all four conditions. It satisfies conditions C1 and C3 because it will always be coarser than \mathbb{C}_{\min} but finer than \mathbb{C}_{\max} . It satisfies conditions C2 and C4 because \mathbb{C} is never updated unless the “valid” subroutine returns True, which implies that C2 and C4 hold.

Finally, \mathbb{C} must be maximal. If a clustering violates C4, then there is no coarser clustering that does not violate C4 since clustering only further reduces information. If a clustering does not violate C2, and a coarser clustering exists that also does not violate C2, then there must exist a way to either merge pairs of clusters or remove variables (namely in reverse topological order) such that each intermediate clustering also does not violate C2. Then, the loop ensures that if a variable could be removed, or two clusters could be merged without violating the conditions, then it will be done. □

Finally, we note that Alg. 3 runs in polynomial time in terms of $n = |\mathbb{V}_L|$ and $|\mathbb{Q}|$, as long as checking identification takes polynomial time. Finding \mathbb{C}_{\min} from Lem. 7 takes polynomial time because finding connected components in an undirected graph takes polynomial time. Checking admissibility simply requires checking cycles, which can also be done in polynomial time. Finding the maximally answerable clusters can also be done in polynomial time, as discussed earlier.

Consider the following example for intuition on choosing the best set of clusters.

Example 16. Consider a setting of annotated image data where

$$\mathbb{V}_L = \{A, T, G, H_1, H_2, O_1, O_2, O_3, O_4, P_1, P_2, P_3, P_4\},$$

Algorithm 3: Choosing intervariable clusters.

Input : Variables \mathbb{V}_L , causal diagram \mathcal{G} , noncausal diagram $\bar{\mathcal{G}}$, queries of interest \mathbb{Q}
Output : Maximally coarse clustering \mathbb{C} such that all criteria are satisfied, or FAIL if none exist.

```

1 Function valid( $\mathbb{C}, \mathcal{G}, \mathbb{Q}$ ):
2   if not admissible( $\mathbb{C}, \mathcal{G}$ ) then
3     return False // C2 violated
4   for  $Q \in \mathbb{Q}$  do
5      $\mathcal{G}_C \leftarrow \text{CDAG}(\mathcal{G}, \mathbb{C})$  // from Def. 9
6     if not ID( $Q, \mathcal{G}_C$ ) then
7       return False // C4 violated
8   return True
9  $\mathbb{C}_{\min} \leftarrow \text{minCluster}(\mathbb{V}_L, \bar{\mathcal{G}})$  // from Lem. 7
10  $\mathbb{C}_{\min} \leftarrow \text{mergeCycles}(\mathbb{C}_{\min}, \mathcal{G})$  // from Lem. 8
11  $\mathbb{C}_{\max} \leftarrow \text{maxCluster}(\mathbb{V}_L, \mathbb{Q})$  // from Lem. 9
12 if not  $\mathbb{C}_{\max}.\text{coarserThan}(\mathbb{C}_{\min})$  then
13   return FAIL // either C1, C2, or C3 violated
14 if not valid( $\mathbb{C}, \mathcal{G}, \mathbb{Q}$ ) then
15   return FAIL // C4 cannot be satisfied
16  $\mathbb{C} \leftarrow \mathbb{C}_{\min}$ 
17 while  $\mathbb{C}$  keeps updating do
18   for  $V \in \mathbb{V}_L$  do
19     if  $V \notin \mathbb{C}$  for all  $\mathbb{C} \in \mathbb{C}_{\max}$  then
20        $\mathbb{C}' \leftarrow \mathbb{C}.\text{remove}(V)$ 
21       if valid( $\mathbb{C}', \mathcal{G}, \mathbb{Q}$ ) then
22          $\mathbb{C} \leftarrow \mathbb{C}'$ 
23   for  $\mathbb{C}_1, \mathbb{C}_2 \in \mathbb{C}$  do
24     for  $\mathbb{C} \in \mathbb{C}_{\max}$  do
25       if  $\mathbb{C}_1, \mathbb{C}_2 \subseteq \mathbb{C}$  then
26          $\mathbb{C}' \leftarrow \text{merge}(\mathbb{C}, \mathbb{C}_1, \mathbb{C}_2)$ 
27         if valid( $\mathbb{C}', \mathcal{G}, \mathbb{Q}$ ) then
28            $\mathbb{C} \leftarrow \mathbb{C}'$ 
29 return  $\mathbb{C}$ 

```

and the causal diagram \mathcal{G} (and noncausal diagram $\bar{\mathcal{G}}$) are given in Fig. 18a. The data is collected from people, and the variables are age (A), gender (G), testosterone level (T), mustache hairs (H_1, H_2), atoms of other mustache hairs (O_1, O_2, O_3, O_4), and pixels of an image of the person (P_1, P_2, P_3, P_4). The queries of interest are $\mathbb{Q} = \{Q_1, Q_2, Q_3\}$, where

$$\begin{aligned} Q_1 &= P(\{p_1, p_2, p_3, p_4\}_{h_1, h_2, o_1, o_2, o_3, o_4}) \\ Q_2 &= P(\{p_1, p_2, p_3, p_4\}_{h_1, h_2, o_1, o_2, o_3, o_4, g}) \\ Q_3 &= P(\{p_1, p_2, p_3, p_4\}_a). \end{aligned}$$

Evidently, the variables depicted in Fig. 18a are too complex, and some intervariable clustering is needed. For starters, there are interactions between the atom variables (O) and the pixel variables (P) that are too low-level to be considered causally. They should be clustered together, as described in Lemma 7. This provides the graph in Fig. 18b, in which all of the O variables are clustered into another hair variable H_3 , and the pixel variables P are clustered into a single image variable I . Formally, this clustering is

$$\begin{aligned} \mathbb{C}_1 &= \{\{A\}, \{T\}, \{G\}, \{H_1\}, \\ &\quad \{H_2\}, H_3 = \{O_1, O_2, O_3, O_4\}, I = \{P_1, P_2, P_3, P_4\}\}. \end{aligned}$$

While this choice of \mathbb{C}_1 is valid in that it satisfies conditions C1-4, there exists a coarser clustering that satisfies the conditions, shown in Fig. 18c. This one is obtained by further clustering H_1, H_2, H_3 into a single “mustache” variable M . Formally, this clustering is

$$\begin{aligned} \mathbb{C}_2 &= \{\{A\}, \{T\}, \{G\}, M = \{H_1, H_2, O_1, O_2, O_3, O_4\}, \\ &\quad I = \{P_1, P_2, P_3, P_4\}\}. \end{aligned}$$

The corresponding queries of $\tau(\mathbb{Q})$, under this choice of \mathbb{C} , would be

$$\begin{aligned} Q_1 &= P(\{p_1, p_2, p_3, p_4\}_{h_1, h_2, o_1, o_2, o_3, o_4}) = P(i_m) \\ Q_2 &= P(\{p_1, p_2, p_3, p_4\}_{h_1, h_2, o_1, o_2, o_3, o_4, g}) = P(i_{m, g}) \\ Q_3 &= P(\{p_1, p_2, p_3, p_4\}_a) = P(i_a). \end{aligned}$$

If for some reason we decide to try to cluster \mathbb{C}_2 further, such as by clustering A, G, M into a single variable X (Fig. 18d), we violate condition C3, the answerability of the queries. How could we distinguish between queries like $P(I_{M=m} = i)$ and $P(I_{A=a} = i)$ if M and A are grouped in the same cluster? For this reason, we cannot proceed any further than \mathbb{C}_2 , and, in fact, \mathbb{C}_2 is actually the maximally answerable clusters from Lemma 9. ■

D.2 Discussion on the Abstract Invariance Condition

Recall the abstract invariance condition (AIC):

Definition 8 (Abstract Invariance Condition (AIC)). Let $\mathcal{M}_L = \langle \mathbf{U}_L, \mathbf{V}_L, \mathcal{F}_L, P(\mathbf{U}_L) \rangle$ be an SCM and $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ be a constructive abstraction function relative to \mathbb{C} and \mathbb{D} . The SCM \mathcal{M}_L is said to satisfy the abstract invariance condition (AIC, for short) with respect to τ if, for all

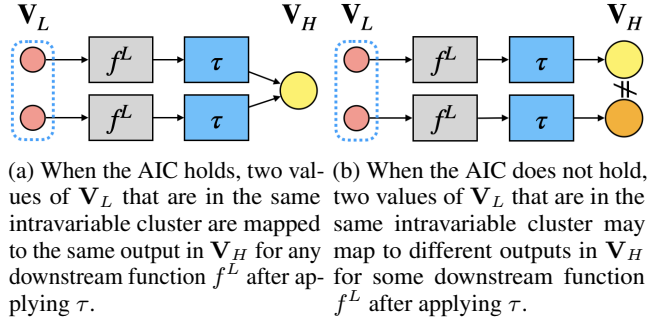


Figure 20: Visualization of the AIC (Def. 8).

$\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{D}_{\mathbf{V}_L}$ such that $\tau(\mathbf{v}_1) = \tau(\mathbf{v}_2), \forall \mathbf{u} \in \mathcal{D}_{\mathbf{U}_L}, \mathbf{C}_i \in \mathbb{C}$, the following holds:

$$\begin{aligned} \tau_{\mathbf{C}_i} \left(\left(f_V^L(\mathbf{pa}_V^{(1)}, \mathbf{u}_V) : V \in \mathbf{C}_i \right) \right) \\ = \tau_{\mathbf{C}_i} \left(\left(f_V^L(\mathbf{pa}_V^{(2)}, \mathbf{u}_V) : V \in \mathbf{C}_i \right) \right), \end{aligned} \quad (33)$$

where $\mathbf{pa}_V^{(1)}$ and $\mathbf{pa}_V^{(2)}$ are the values corresponding to \mathbf{v}_1 and \mathbf{v}_2 . Then, $\widetilde{\mathbf{pa}}_V$ is used to denote any arbitrary value s.t. $\tau(\widetilde{\mathbf{pa}}_V) = \tau(\mathbf{pa}_V^{(1)}) = \tau(\mathbf{pa}_V^{(2)})$. ■

The AIC intuitively ensures that there is no loss of information when clustering two values together in the intravariabile clusters \mathbb{D} . Specifically, it states that clustering two values together is safe whenever these values are interchangeable with respect to the behavior of downstream functions. That is, the functions are *invariant* to changes between values in the same intravariabile cluster. This is illustrated in Fig. 20. Examples 4 and 5 from Sec. 2 illustrate the subtleties of the AIC both when it holds and when it does not hold.

If the AIC is violated, then two functionally different values were placed in the same cluster. Hence, assuming that the AIC holds is reasonable provided that the intravariabile clusters were chosen reasonably. For example, in a study where a drug X is reasonably expected to affect blood pressure levels Y , it would not make sense to cluster the values $X = 0$ (drug not taken) and $X = 1$ (drug taken) together since, by design, they have different downstream effects on Y and cannot be treated similarly.

On the other hand, in an image recognition setting, perhaps it does not matter for the task if the image is scaled or rotated, so these invariances can be modeled in the intravariabile clustering. For example, if one image is a scaled version of another but they are otherwise the same, the two images might be placed in the same intravariabile clustering without violating the AIC. Sec. 4 and Appendix D.3 discuss this phenomenon in more detail.

If a data scientist finds the AIC too strict for her settings, one option is to revise the intravariabile clusters, perhaps making them less coarse. Otherwise, the abstraction may be too strong and some important information is lost.

Still, there are alternatives if weaker assumptions are desired. The AIC ensures that an \mathcal{L}_3 - τ consistent model exists, as demonstrated by Alg. 1, but it is not a necessary condition in applications that do not require full \mathcal{L}_3 - τ consistency. For

example, if one only works on the interventional level (\mathcal{L}_2), one may not necessarily care if consistencies between \mathcal{L}_3 quantities are lost.

In the strictest form, making as minimal assumptions as needed for the task, a query-specific AIC can be assumed. That is,

Definition 25. Let \mathbb{Q} be a set of queries of the form $P^{\mathcal{M}_H}(\mathbf{y}_{H,1[\mathbf{x}_{H,1}]}, \mathbf{y}_{H,2[\mathbf{x}_{H,2}]}, \dots)$. For all $Q \in \mathbb{Q}$, we have

$$\begin{aligned} & \sum_{\forall i_{\mathbf{y}_{L,i}} \in \mathcal{D}_{\mathbf{y}_{L,i}} : \tau(\mathbf{y}_{L,i}) = \mathbf{y}_{H,i}} P^{\mathcal{M}_L}(\mathbf{y}_{L,1[\mathbf{x}_{L,1}]}, \mathbf{y}_{L,2[\mathbf{x}_{L,2}]}, \dots) \\ &= \sum_{\forall i_{\mathbf{y}_{L,i}} \in \mathcal{D}_{\mathbf{y}_{L,i}} : \tau(\mathbf{y}_{L,i}) = \mathbf{y}_{H,i}} P^{\mathcal{M}_L}(\mathbf{y}_{L,1[\mathbf{x}'_{L,1}]}, \mathbf{y}_{L,2[\mathbf{x}'_{L,2}]}, \dots), \end{aligned}$$

for all $\mathbf{x}_{L,i}, \mathbf{x}'_{L,i}$ such that $\tau(\mathbf{x}_{L,i}) = \tau(\mathbf{x}'_{L,i}) = \mathbf{x}_{H,i}$. ■

In words, τ satisfies Def. 25 if and only if there exists \mathcal{M}_H that is Q - τ consistent with \mathcal{M}_L for all $Q \in \mathbb{Q}$, since all such variations of Q have the same probability. Still, it may be more useful to have a general criterion that is not query-dependent.

It turns out that the study of weaker forms of the AIC is not new and is first discussed in Chalupka, Perona, and Eberhardt (2015). Broadly, the paper studies a similar abstraction setting in which causal macro-variables are learned from observational data from the low-level micro-variables. Similar to the intravariability learning problem posed in Sec. 4, the paper aims to learn intravariability clusters of a low-level high-dimensional variable I (e.g., an image), which causes some other variable T , possibly with unobserved confounding. However, while the results of Sec. 4 focus on satisfying the AIC in Def. 8, Chalupka, Perona, and Eberhardt (2015) focuses on satisfying weaker conditions that are verifiable through data.

More specifically, consider the setting proposed in Chalupka, Perona, and Eberhardt (2015) illustrated by the causal diagram in Fig. 21. Note that Eq. 33 in Def. 8 requires that

$$\tau(f_T(i_1, \mathbf{u}_T)) = \tau(f_T(i_2, \mathbf{u}_T)), \quad (102)$$

for any $i_1, i_2 \in \mathcal{D}_I$ that are chosen to be clustered in the same intravariability cluster. This condition is a requirement on the SCM function f_T , which, as demonstrated by Lem. 5, implies that one cannot cluster together any two values i_1, i_2 without further information about f_T , which may be difficult to acquire due to the unobserved nature of the SCM. Instead, the paper proposes two other properties. Notably, two values i_1 and i_2 are part of the same *observational partition* (Chalupka, Perona, and Eberhardt 2015, Def. 1) if

$$P(t | i_1) = P(t | i_2), \quad (103)$$

and they are part of the same *causal partition* (Chalupka, Perona, and Eberhardt 2015, Def. 3) if

$$P(t_{i_1}) = P(t_{i_2}). \quad (104)$$

Values of \mathcal{D}_I can easily be clustered according to Eq. 103 using the observational data. However, Eq. 104 may be a more desirable property if the goal of downstream tasks is to perform causal inferences from I to T .

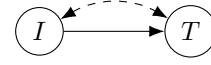


Figure 21: Problem setting in Chalupka, Perona, and Eberhardt (2015).

It turns out that a clustering learned according to Eq. 103 typically also satisfies Eq. 104 according to a major proven result of the paper, the *Causal Coarsening Theorem* (Chalupka, Perona, and Eberhardt 2015, Thm. 5), which states (in informal terms) that the maximally coarse clustering that satisfies Eq. 104 is almost always a coarser clustering of the maximally coarse clustering that satisfies Eq. 103. This essentially implies that, under certain faithfulness conditions, one can achieve the property of Eq. 104 simply as a consequence of achieving Eq. 103, which can be accomplished with observational data alone.

The concepts of the partitions described by Eqs. 103 and 104 can be generalized to the general SCM/graph case. Suppose we are given the inter/intravariability clusters \mathbb{C} and \mathbb{D} . Consider a variant of the AIC applied on conditional probabilities:

Definition 26 (Conditional Abstract Invariance Condition). Let $P(\mathbf{V}_L)$ be an observational distribution over \mathbf{V}_L , and let $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ be a constructive abstraction function relative to \mathbb{C} and \mathbb{D} . $P(\mathbf{V}_L)$ is said to satisfy the conditional abstract invariance condition (conditional AIC, for short) w.r.t. τ if, for all \mathbf{X}_L where \mathbf{X}_L is a union of clusters of \mathbb{C} , all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{D}_{\mathbf{X}_L}$ such that $\tau(\mathbf{x}_1) = \tau(\mathbf{x}_2)$, and all $\mathbf{v}_H \in \mathcal{D}_{\mathbf{V}_H}$, we have

$$\sum_{\mathbf{v}_L \in \mathcal{D}_{\mathbf{V}_L} : \tau(\mathbf{v}_L) = \mathbf{v}_H} P(\mathbf{v}_L | \mathbf{x}_1) = \sum_{\mathbf{v}_L \in \mathcal{D}_{\mathbf{V}_L} : \tau(\mathbf{v}_L) = \mathbf{v}_H} P(\mathbf{v}_L | \mathbf{x}_2). \quad (105)$$

That is, two values can be clustered together as long as their conditional probabilities do not change in the space of \mathbf{V}_H . Note that this condition is entirely focused on the observational distributions $P(\mathbf{V}_L)$ and $P(\mathbf{V}_H)$ and make no requirements over any distributions from the higher layers, \mathcal{L}_2 and \mathcal{L}_3 .

Now consider an interventional variant:

Definition 27 (Interventional Abstract Invariance Condition). Let $\mathcal{L}_2(\mathcal{M}_L)$ be a collection of interventional distributions over \mathbf{V}_L , and let $\tau : \mathcal{D}_{\mathbf{V}_L} \rightarrow \mathcal{D}_{\mathbf{V}_H}$ be a constructive abstraction function relative to \mathbb{C} and \mathbb{D} . We say that $\mathcal{L}_2(\mathcal{M}_L)$ satisfies the interventional abstract invariance condition (interventional AIC, for short) w.r.t. τ if, for all \mathbf{X}_L where \mathbf{X}_L is a union of clusters of \mathbb{C} , all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{D}_{\mathbf{X}_L}$ such that $\tau(\mathbf{x}_1) = \tau(\mathbf{x}_2)$, and all $\mathbf{v}_H \in \mathcal{D}_{\mathbf{V}_H}$, we have

$$\sum_{\mathbf{v}_L \in \mathcal{D}_{\mathbf{V}_L} : \tau(\mathbf{v}_L) = \mathbf{v}_H} P(\mathbf{v}_L | \mathbf{x}_1) = \sum_{\mathbf{v}_L \in \mathcal{D}_{\mathbf{V}_L} : \tau(\mathbf{v}_L) = \mathbf{v}_H} P(\mathbf{v}_L | \mathbf{x}_2). \quad (106)$$

This states that two values can be clustered together as long as the corresponding interventional distributions do not change in the space of \mathbf{V}_H .

These two definitions may be useful when looking to find abstractions while working on lower layers of the PCH. When given the C-DAG \mathcal{G}_C , these can be refined to specific queries based on parent-child relationships in the graph. Although not shown formally here, it is likely that a more general variant of the Causal Coarsening Theorem could be proven for these two definitions. Nonetheless, using these definitions in place of the full AIC from Def. 8 has its limitations; for example, it may no longer be possible to perform counterfactual inferences on the high-level model. We encourage further research on this topic, comparing the relationships between variations of the AIC.

D.3 Representation Learning with Intravari-able Clusters

Recall the following proposition from Sec. 4:

Proposition 5. *Consider a low level SCM \mathcal{M}_L and constructive abstraction function τ w.r.t. clusters \mathbb{C} and \mathbb{D} . \mathcal{M}_L is guaranteed to satisfy the AIC w.r.t. τ if and only if $\mathbb{D}_{C_i} = \{\{c_i\} : c_i \in \mathcal{D}_{C_i}\}$ for all $C_i \in \mathbb{C}$. ■*

When learning intravari-able clusters, Prop. 5 states that without extra information, each value needs to be put in its own cluster. However, if extra information is available for use, this can be leveraged to improve the choice of \mathbb{D} . In practice, this is common, and is generally provided in the form of invariance assumptions. As a intuitive example, suppose that X is an image of an animal, Y is the corresponding label, and $f_Y(x, u_Y)$ describes the mechanism that humans use to classify X (here, u_Y is the exogenous noise). Suppose that $g(x, \phi)$ is a function that rotates x by ϕ degrees. It is the case that the classification of animals is rotationally invariant, that is, $f_Y(g(x, \phi), u_Y) = f_Y(x, u_Y)$ for all x, ϕ, u_Y . Then, in fact, \mathbb{D} can be defined such that values of x_1, x_2 are clustered together if $g(x_1, \phi) = x_2$, as shown by the following result.

Proposition 12. *For each $C_i \in \mathbb{C}$, suppose that there exists function $g_{C_i} : \mathcal{D}_{C_i} \times \mathcal{D}_{\phi_{C_i}} \rightarrow \mathcal{D}_{C_i}$ (with parameters ϕ_{C_i}) such that*

$$f_V^L(c_i, u_V) = f_V^L(g_{C_i}(c_i, \phi_{C_i}), u_V) \quad (107)$$

for all V that are children of C_i and all $\phi_{C_i} \in \mathcal{D}_{\phi_{C_i}}$. Then, \mathcal{M}_L satisfies the AIC w.r.t. any constructive abstraction function with \mathbb{D} such that $c_i^{(1)}, c_i^{(2)}$ being in the same intravari-able cluster implies that $g_{C_i}(c_i^{(1)}, \phi_{C_i}) = c_i^{(2)}$ for some ϕ_{C_i} . ■

Proof. Consider the proposed set of intravari-able clusters \mathbb{D} . Let $v_1, v_2 \in \mathcal{D}_{V_L}$ be two values such that $\tau(v_1) = \tau(v_2)$. The goal is to show that for all $u \in \mathcal{D}_{U_L}$ and $C_k \in \mathbb{C}$, we have

$$\begin{aligned} & \tau\left(\left(f_V^L(\mathbf{pa}_V^{(1)}, u_V) : V \in C_k\right)\right) \\ &= \tau\left(\left(f_V^L(\mathbf{pa}_V^{(2)}, u_V) : V \in C_k\right)\right), \end{aligned} \quad (108)$$

where $\mathbf{pa}_V^{(1)}$ and $\mathbf{pa}_V^{(2)}$ are the values corresponding to v_1 and v_2 respectively. If $\tau(v_1) = \tau(v_2)$, that implies that for all $C_i \in \mathbb{C}$ (where $c_i^{(1)}$ and $c_i^{(2)}$ correspond to the values of

v_1 and v_2 respectively), $c_i^{(1)}$ and $c_i^{(2)}$ must be in the same intravari-able cluster in \mathbb{D}_{C_i} (this includes the case where $c_i^{(1)} = c_i^{(2)}$). By construction of \mathbb{D} , this is only possible if $g_{C_i}(c_i^{(1)}, \phi_{C_i}) = c_i^{(2)}$ for some ϕ_{C_i} .

For $g_V = \{g_{C_i} : C_i \cap \mathbf{Pa}_V \neq \emptyset\}$, denote $g_V(\mathbf{pa}_V^{(1)}, \phi)$ as the collection of outputs of $g_{C_i}(c_i^{(1)}, \phi_{C_i})$ for the functions of $g_{C_i} \in g_V$, which, as established earlier, is equal to $\mathbf{pa}_V^{(2)}$. Then we have

$$\begin{aligned} & \tau\left(\left(f_V^L(\mathbf{pa}_V^{(1)}, u_V) : V \in C_k\right)\right) \\ &= \tau\left(\left(f_V^L(g_V(\mathbf{pa}_V^{(1)}, \phi), u_V) : V \in C_k\right)\right) \\ &= \tau\left(\left(f_V^L(\mathbf{pa}_V^{(2)}, u_V) : V \in C_k\right)\right), \end{aligned}$$

completing the proof. □

For intuition, consider the following simple example.

Example 17. Consider a situation where Y is some binary variable that is a (noisy) bitwise AND of two other binary variables X_1 and X_2 . For example, perhaps Y denotes whether a law is enacted, and X_1 and X_2 denotes the votes of the two branches of government, like in Ex. 11. Formally, let \mathcal{M}_L be defined as follows:

$$\begin{aligned} \mathbf{U}_L &= \{U_{X_1}, U_{X_2}, U_Y\}, \text{ all binary} \\ \mathbf{V}_L &= \{X_1, X_2, Y\}, \text{ all binary} \\ \mathcal{F}_L &= \begin{cases} X_1 \leftarrow f_{X_1}^L(u_{X_1}) & = u_{X_1} \\ X_2 \leftarrow f_{X_2}^L(u_{X_2}) & = u_{X_2} \\ Y \leftarrow f_Y^L(x_1, x_2, u_Y) & = (x_1 \wedge x_2) \oplus u_Y \end{cases} \\ P(\mathbf{U}_L) &= \begin{cases} P(U_{X_1} = 1) & = P(U_{X_2} = 1) = 0.5 \\ P(U_Y = 1) & = 0.1 \end{cases} \end{aligned}$$

Now suppose we want to create an abstraction of \mathcal{M}_L using the intervari-able clusters $\mathbb{C} = \{C_1 = \{X_1, X_2\}, C_2 = \{Y\}\}$. The domain \mathcal{D}_{C_1} has four values as X_1 and X_2 can both be either 0 or 1. However, as emphasized by Prop. 5, the only intravari-able clustering we can choose without additional information is the one where each value is in its own cluster. In other words, $\mathbb{D}_{C_1} = \{\{(X_1 = 0, X_2 = 0)\}, \{(X_1 = 0, X_2 = 1)\}, \{(X_1 = 1, X_2 = 0)\}, \{(X_1 = 1, X_2 = 1)\}\}$.

However, suppose we are given additional information that f_Y is permutation invariant to its endogenous inputs, i.e.,

$$f_Y(x_1, x_2, u_Y) = f_Y(x_2, x_1, u_Y). \quad (109)$$

In the notation of Prop. 12, we can define ϕ_{C_1} to be a binary variable such that 0 means original order and 1 means reversed. Then we can define $g(x_1, x_2, \phi_{C_1}) = (x_1, x_2)$ if $\phi_{C_1} = 0$ or (x_2, x_1) if $\phi_{C_1} = 1$. Then, this implies that $f_Y(x_1, x_2, u_Y) = f_Y(g(x_1, x_2, \phi_{C_1}), u_Y)$ for any choice of ϕ_{C_1} .

By Prop. 12, this implies that $(X_1 = 0, X_2 = 1)$ and $(X_1 = 1, X_2 = 0)$ can be placed in the same intravari-able cluster without violating the AIC. Indeed, the function

$f_Y^L(x_1, x_2, u_Y) = (x_1 \wedge x_2) \oplus u_Y$ does not change when x_1 and x_2 are swapped, so we have

$$\begin{aligned} & \tau \left((f_Y^L(X_1 = 0, X_2 = 1, \mathbf{u}_V) : V \in \mathbf{C}_k) \right) \\ &= \tau \left((f_Y^L(X_1 = 1, X_2 = 0, \mathbf{u}_V) : V \in \mathbf{C}_k) \right), \end{aligned}$$

confirming that the AIC still holds when $(X_1 = 0, X_2 = 1)$ and $(X_1 = 1, X_2 = 0)$ can be placed in the same intravari-able cluster. ■

In practice, this invariance information can be incorporated in the process of learning the intravari-able clusters while training $\hat{\tau}$ in an RNCM, as described in Def. 11. To take into account Prop. 12, the training objective could include a term to enforce the invariance specified by $g_{\mathbf{C}_i}$ (e.g. through a penalty on $\tau_{\mathbf{C}_i}(\mathbf{c}_i, \boldsymbol{\theta}_{\mathbf{C}_i}) - \tau_{\mathbf{C}_i}(g_{\mathbf{C}_i}(\mathbf{c}_i, \phi_{\mathbf{C}_i}), \boldsymbol{\theta}_{\mathbf{C}_i})$).

In fact, it turns out that the idea of intravari-able clusters works in tandem with techniques in deep representation learning when it comes to incorporating invariances in the data. In the case of image data for example, many works in com-puter vision have already leveraged general patterns found in images (e.g. rotations, crops, flips, etc. do not affect clas-sification) to achieve faster training with less data (Shorten and Khoshgoftaar 2019). Permutation invariance concepts (Zaheer et al. 2017; Murphy et al. 2019) have been used as pooling functions for convolutional neural networks (LeCun and Bengio 1998). Existing frameworks for representation learning such as through contrastive methods (Chen et al. 2020) can be used for learning τ .

E Additional Examples

This section contains examples to improve the clarity of concepts in the paper.

E.1 Sec. 2 Examples

We provide more examples of the concepts in Sec. 2, namely, abstractions constructed through inter/intravari-able clusters. In addition to Example 1 within the section, we provide a more involved example below.

Example 18. Suppose an economist is studying the effects of implementing a new type of government policy on recession prevention. The economist records data on several variables of interest: whether the policy is implemented (X); whether the policy is lobbied (W); economic spending in terms of consumption (C), investment (I), government spending (G), imports (M), and exports (E); and whether or not there is a recession (Y). Out of these variables, Z , X , and Y are binary, and C , I , G , M , and E are numerical values representing how much the spending has .changed relative to the previous year (in billions of dollars).

Suppose the true SCM $\mathcal{M}^* = \mathcal{M}_L$ is defined as follows:

$$\mathbf{U}_L = \{U_W, U_X, U_{XY}, U_C, U_I, U_G, U_M, U_E, U_{IG}, U_{CM}, U_{EM}\}$$

$$\mathbf{V}_L = \{W, X, C, I, G, M, E, Y\}$$

$$\mathcal{F}_L = \{$$

$$W \leftarrow f_W(u_W) = u_W$$

$$X \leftarrow f_X(w, u_X, u_{XY}) = w \oplus u_X \oplus u_{XY}$$

$$C \leftarrow f_C(x, g, u_C) = 5x - 0.2g + u_C + u_{CM}$$

$$I \leftarrow f_I(x, u_I) = -5x + u_I + u_{IG}$$

$$G \leftarrow f_G(x, u_G) = 10x + u_G + u_{IG}$$

$$E \leftarrow f_E(u_E) = u_E + u_{EM}$$

$$M \leftarrow f_M(c, u_M) = 0.2c + u_M + u_{CM} - u_{EM}$$

$$Y \leftarrow f_Y(c, i, g, e, m, u_{XY}) =$$

$$\mathbf{1}\{c + i + g + e - m \leq 0\} \oplus u_{XY}$$

$$P(\mathbf{U}_L) = \begin{cases} U_W, U_X & \sim \text{Bernoulli}(0.5) \\ U_{XY} & \sim \text{Bernoulli}(0.1) \\ U_C, U_I, U_G, U_M, U_E & \sim N(0, 10) \\ U_{IG}, U_{CM}, U_{EM} & \sim N(0, 2) \end{cases}$$

To summarize, the policy has some impact on the con-sumption, investment, and government spending. Whether or not there is a recession depends on imports and exports in ad-dition to all of these factors. Indeed, despite the fact that this is a toy example, the SCM \mathcal{M}^* is already quite complex to describe with this many variables and functions, and perhaps this level of detail is not needed to achieve the inference we desire. We will see how utilizing abstractions can help with this.

We first note that perhaps \mathbf{V}_L contains too many micro-level variables that can be summarized with a smaller set of abstract higher-level variables. Further, perhaps using real-valued variables is overly complex, and the same phenomena can be described without loss of generality using a lower-dimensional space. For entertaining these considerations, we utilize the idea of intervariable and intravari-able clusters.

Suppose we are only interested in the causal effect of the policy X on recession Y . We could study the same phe-nomenon under a simpler set of variables. To do so, we can cluster variables of \mathbf{V}_L to form a new set of macro-level variables. First, note that although we may have data on C , I , G , E , and M , it may make more sense to simply cluster them together and consider them as one variable (e.g. GDP). Further, perhaps we may decide that W is not relevant to the analysis and exclude it from the study. We leave these clustering decisions at the discretion of the data scientist (e.g. the economist in this example).

By the definition of intervariable clusters from Def. 5, we can choose clusters $\mathbb{C} = \{\mathbf{C}_1 = \{X\}, \mathbf{C}_2 = \{C, I, G, E, M\}, \mathbf{C}_3 = \{Y\}\}$. We leave X and Y in their own clusters while grouping all of C , I , G , E , M into one cluster \mathbf{C}_2 . W is excluded from all of the clusters and is effectively projected out of the system. We can then treat \mathbf{C}_1 , \mathbf{C}_2 , and \mathbf{C}_3 as our new variables. Let us relabel them X , Z , and Y respectively, where Z . We can define our higher level variables as $\mathbf{V}_H = \{X, Z, Y\}$.

Further, it is not immediately clear what the domain of Z is. Certainly, it could be left as a tuple $(C, I, G, E, M) \in \mathbb{R}^5$. However, this would be an overly complex representation of Z , and we do not need to retain all of the joint information of C, I, G, E, M . In other words, \mathcal{D}_Z does not have to be equal to $\mathcal{D}_{\mathbf{C}_2}$ and can be represented more compactly. Instead, we may choose to define Z as the GDP, or $Z = C + I + G - E - M$. This can be described through the use of intravariabile clusters.

By the definition of intravariabile clusters from Def. 5, we must choose $\mathbb{D} = \{\mathbb{D}_{\mathbf{C}_1}, \mathbb{D}_{\mathbf{C}_2}, \mathbb{D}_{\mathbf{C}_3}\}$, where $\mathbb{D}_{\mathbf{C}_i}$ is a partitioning of the domain $\mathcal{D}_{\mathbf{C}_i}$. Since X and Y are already binary variables, we will not be able to group their values together, so we simply define $\mathbb{D}_{\mathbf{C}_1}$ and $\mathbb{D}_{\mathbf{C}_2}$ as $\{\{0\}, \{1\}\}$. For \mathbf{C}_2 however, we must cluster values of (C, I, G, E, M) such that $(c_1, i_1, g_1, e_1, m_1), (c_2, i_2, g_2, e_2, m_2) \in \mathcal{D}_{\mathbf{C}_2}$ are in the same cluster if and only if $c_1 + i_1 + g_1 + e_1 - m_1 = c_2 + i_2 + g_2 + e_2 - m_2$. In other words, we can, for example, define $\mathbb{D}_{\mathbf{C}_2} = \{\mathcal{D}_{\mathbf{C}_2}^j : (c, i, g, e, m) \in \mathcal{D}_{\mathbf{C}_2}, c + i + g + e - m = j\}$.

Now, with the intravariabile clusters \mathbb{D} defined, we can choose the domains of \mathbf{V}_H as simply their corresponding clusters in \mathbb{D} . That is, each value of Z corresponds to some $\mathcal{D}_{\mathbf{C}_2}^j$. In fact, we can simply set $Z = j$, where $j = C + I + G + E - M$, which intuitively corresponds to the idea that Z represents the annual change in GDP. Note that the domain of Z , \mathcal{D}_Z becomes smaller than $\mathcal{D}_{\mathbf{C}_2}$ in some sense. First, it is lower dimensional (\mathbb{R} instead of \mathbb{R}^5), and second, there are clearly values of (C, I, G, E, M) that are mapped to the same value of Z . For example, $(C = 1, I = 2, G = 3, E = 4, M = 2)$ maps to the same value as $(C = 3, I = 2, G = 1, E = 5, M = 3)$ because $1 + 2 + 3 + 4 - 2 = 3 + 2 + 1 + 5 - 3 = 8$.

However, perhaps this level of dimensionality reduction is insufficient. In terms of cardinality, \mathcal{D}_Z is the same size as $\mathcal{D}_{\mathbf{C}_2}$ as the cardinality of \mathbb{R} and \mathbb{R}^5 are the same. Perhaps the domain of Z could be compressed further. Indeed, we could turn Z into a binary variable by defining $\mathbb{D}_{\mathbf{C}_2} = \{\mathcal{D}_{\mathbf{C}_2}^0, \mathcal{D}_{\mathbf{C}_2}^1\}$, where $(c, i, g, e, m) \in \mathcal{D}_{\mathbf{C}_2}^0$ if $c + i + g + e - m > 0$ or it is in $\mathcal{D}_{\mathbf{C}_2}^1$ otherwise. In other words, we drop all information about the annual change in GDP except whether it is positive or negative. We will use this clustering for the rest of the examples, and we will see in later examples why this clustering is allowed and makes sense. Sec. 4 expands on the general discussion of which intravariabile cluster choices are allowed.

With the idea of inter/intravariabile clusters, the concept of constructive abstraction functions can be established as in Def. 6. From Example 18, the function τ constructed from \mathbb{C} and \mathbb{D} is clear. For example, consider the value $\mathbf{v}_L = (W = 0, X = 1, C = 2.6, I = -1.2, G = 10.2, E = 0.4, M = 1.2, Y = 0)$. We can compute $\mathbf{v}_H = \tau(\mathbf{v}_L)$ and find that $\mathbf{v}_H = (X = 1, Z = 0, Y = 0)$. In this case, $\tau = (\tau_{\mathbf{C}_1}, \tau_{\mathbf{C}_2}, \tau_{\mathbf{C}_3})$, where $\tau_{\mathbf{C}_1}$ and $\tau_{\mathbf{C}_3}$ are the identity function on X and Y respectively, and $\tau_{\mathbf{C}_2}(C = 2.6, I = -1.2, G = 10.2, E = 0.4, M = 1.2) = \mathbf{1}\{2.6 - 1.2 + 10.2 + 0.4 - 1.2 \leq 0\} = 0$. That is, for \mathbf{C}_2 , the intravariabile clusters are defined such that values of (c, i, g, e, m) are divided into two categories depending on whether $c + i + g + e - m \leq 0$,

which we can arbitrarily choose as the binary values 0 and 1. This provides a mapping from every value of $\mathcal{D}_{\mathbf{V}_L}$ to some value of $\mathcal{D}_{\mathbf{V}_H}$. ■

Note that τ alone does not define an abstraction. While τ provides a well-defined mapping from \mathbf{V}_L to \mathbf{V}_H , not every SCM defined over \mathbf{V}_H can be considered an abstraction of \mathcal{M}^* . Consider the following example.

Example 19. Consider the SCM \mathcal{M}_H defined as follows:

$$\mathcal{M}_H = \begin{cases} \mathbf{U}_H & = \{U_X, U_Z, U_Y\} \\ \mathbf{V}_H & = \{X, Z, Y\} \\ \mathcal{F}_H & = \begin{cases} X \leftarrow f_X(u_X) & = u_X \\ Z \leftarrow f_Z(u_Z) & = u_Z \\ Y \leftarrow f_Y(u_Y) & = u_Y \end{cases} \\ P(\mathbf{U}_H) & = U_X, U_Z, U_Y \sim \text{Bernoulli}(0.5) \end{cases}$$

One could argue that \mathcal{M}_H is defined over \mathbf{V}_H , constructed via τ from Example 18. However, it is trivial to see that \mathcal{M}_H retains none of the meaning of the variables of \mathbf{V}_H intended by τ and is clearly oblivious of \mathcal{M}_L . After all, there is not even any causal relationship between the variables defined in \mathcal{M}_H . Intuitively, \mathcal{M}_H is clearly not an abstraction of \mathcal{M}^* , and this is reinforced by the fact that \mathcal{M}_H does not match common definitions of abstractions, such as Def. 15.

However, assuming that we do not have access to \mathcal{M}_H or \mathcal{M}^* , and can only observe them through their distributions of the PCH, how could we tell that \mathcal{M}_H is not an abstraction of \mathcal{M}^* ? Certainly, there are inconsistencies in the distribution too. For example, note that

$$\begin{aligned} P^{\mathcal{M}^*}(C + I + G + E - M > 0 \mid X = 1) &= P(5 - 0.2g + u_C + u_{CM} - 5 + u_I + u_{IG} + 10 + u_G \\ &+ u_{IG} + u_E + u_{EM} - 0.2c - u_M - u_{CM} - u_{EM} > 0) \\ &= P(5 - 0.2(10 + u_G + u_{IG}) + u_C + u_{CM} - 5 + u_I \\ &+ u_{IG} + 10 + u_G + u_{IG} + u_E + u_{EM} \\ &- 0.2(5 - 0.2(10 + u_G + u_{IG}) + u_C + u_{CM}) \\ &- u_M - u_{CM} - u_{EM} > 0) \\ &= P(7.4 + 0.84u_G + 1.8u_{IG} + 0.8u_C - 0.2u_{CM} \\ &+ 1.04u_I + u_E - u_M > 0) \\ &= P(N(0, 7.056) + N(0, 6.48) + N(0, 6.4) + N(0, 0.08) \\ &+ N(0, 10.816) + N(0, 10) + N(0, 10) > -7.4) \\ &= P(N(0, 50.832) > -7.4) \approx 0.85. \end{aligned}$$

On the other hand, we see that $P^{\mathcal{M}_H}(Z = 0 \mid X = 1) = P^{\mathcal{M}_H}(Z = 0) = P(U_Z = 0) = 0.5$. We would expect that if \mathcal{M}_H were an abstraction of \mathcal{M}^* , then $P^{\mathcal{M}_H}(Z = 0 \mid X = 1)$ should match $P^{\mathcal{M}^*}(C + I + G + E - M > 0 \mid X = 1)$, which is evidently not the case.

The intuition behind this connection is that τ directly maps cases of $C + I + G + E - M > 0$ to cases of $Z = 0$, so their corresponding probabilities should be the same. If $Q = P^{\mathcal{M}^*}(C + I + G + E - M > 0 \mid X = 1)$, we would say that \mathcal{M}_H is not Q - τ consistent with \mathcal{M}^* , via Def. 7. ■

Not all functions mapping between two spaces of variables are constructive abstraction functions. Some abstraction functions, while not being constructive, may still have qualities

that are expected from abstractions. However, the lack of structure in these functions make certain features of the abstraction models difficult to define, such as the concept of Q - τ consistency. The following example describes a few such cases.

Example 20. To start, we note that when fixing the intervariable and intravariabile clusters \mathbb{C} and \mathbb{D} , the corresponding constructive abstraction function τ is deterministic, and the corresponding \mathbf{V}_H is fixed, up to a bijective mapping of its domain. For example, suppose $\mathbb{C} = \{A, B, C\}$, and $\mathbb{D}_A = \{a_1, a_2\}$, $\mathbb{D}_B = \{b_1, b_2\}$, $\mathbb{D}_C = \{c_1, c_2\}$. However, let $\mathbf{V}_H = \{X, Y, Z, W\}$, where $\mathcal{D}_X = \mathcal{D}_Y = \mathcal{D}_Z = \{0, 1\}$ and $\mathcal{D}_W = \{0, 1, 2\}$. While a function could be constructed from \mathbb{D} and \mathbb{C} , it could not be a valid mapping to \mathbf{V}_H . First of all, there are four variables in \mathbf{V}_H but only three clusters in \mathbb{C} , so the mapping could not be one-to-one.

Even if X were removed, all of the intervariable clusters have values clustered into two sets, but W is a ternary variable, so a one-to-one mapping is not possible with any of the intervariable clusters. However, if W were removed from \mathbf{V}_H , then τ could be a valid constructive abstraction mapping between \mathbf{V}_L and \mathbf{V}_H , since the other three variables are binary. The number of variables of \mathbf{V}_H and their domain sizes match the number of clusters of \mathbb{C} and the number within each cluster of \mathbb{D} , respectively. In fact, at least without considering the model distributions, any cluster of \mathbb{C} could map to any of the variables of $\mathbf{V}_H \setminus \{W\}$ as they are isomorphic.

In other words, τ fails to be a constructive tau abstraction for any choice of \mathbf{V}_H that is incompatible with \mathbb{C} and \mathbb{D} . Intuitively, this means that when given \mathbb{C} and \mathbb{D} , the function τ and high level variables \mathbf{V}_H are already fixed, which is how the rest of this paper is framed. However, there can exist cases with some function τ mapping from \mathbf{V}_L to \mathbf{V}_H such that τ cannot be considered a constructive abstraction function for any choice of \mathbb{C} and \mathbb{D} . Some cases may even appear to be valid abstractions intuitively.

Consider an example of a company board setting. Suppose in a company, the board consists of the CEO, Alice (A), and two vice presidents, Bob (B) and Charlie (C). When voting on company policies, each board member can choose to vote for (+1), vote against (-1), or abstain (0). In other words, A, B , and C are ternary variables with domain $\{-1, 0, +1\}$. While data can be collected on each of the members' voting behaviors ($\mathbf{V}_L = \{A, B, C\}$), it may be more sensible to aggregate the votes into a more useful quantity. Suppose two high level variables are computed ($\mathbf{V}_H = \{X, Z\}$) through some abstraction function $\tau = (\tau_X, \tau_Z)$, defined as follows.

$$\tau_X(a, b, c) = \begin{cases} +1 & a + b + c > 0 \\ 0 & a + b + c = 0 \\ -1 & a + b + c < 0 \end{cases}$$

$$\tau_Z(a, b, c) = \begin{cases} a & a \neq 0 \\ +1 & a = 0, b + c > 0 \\ 0 & a = 0, b + c = 0 \\ -1 & a = 0, b + c < 0 \end{cases}$$

In words, X is the aggregate vote that is simply the majority vote of all three members. On the other hand Z is an ag-

gregate vote that prioritizes Alice's vote, as she is CEO. If she abstains, then it is an aggregate of the votes of the vice presidents. For example, $\tau(A = +1, B = -1, C = -1) = (X = -1, Z = +1)$.

Perhaps in company matters, it is more useful to use the variables X and Z over the individual votes of the board members. However, τ cannot be considered a constructive abstraction function for any choice of clusters \mathbb{C} and \mathbb{D} . This is because both X and Z change values depending on all of the values of A, B , and C , so the variables of \mathbf{V}_L cannot be cleanly separated into two different clusters.

Even so, this choice of τ seems like it could result in a valid abstraction. The difficulty lies with the analysis of causal quantities after τ is fixed. Without the notion of clusters, the definition of Q - τ consistency fails to work, so it is no longer clear what causal quantities correspond to what. For example, what would an intervention of $A = +1$ imply on the \mathbf{V}_H level? Or, what would an intervention on $Z = +1$ imply on the \mathbf{V}_L level? The answer might change depending on the setting, or there may not even be an answer that makes sense at all. This is why prior works like Rubenstein et al. (2017) define the mapping between interventions separately from τ , which maps the variables.

In this particular case, provided that queries do not require separation of X and Z (e.g. queries like $P(X_{Z=z} = x)$ are not needed), then it may be sensible to cluster A, B , and C together into one intervariable cluster and then simply have one variable in \mathbf{V}_H used for downstream tasks. This allows for the theory in this paper to be applied, significantly reducing the complexity of defining the abstractions. ■

See the following for a negative example of an abstraction of the drug example in Ex. 3.

Example 21. Continuing Example 3, consider the following SCM \mathcal{M}_H .

$$\mathbf{U}_H = \{U_X, U_Y\}$$

$$\mathbf{V}_H = \{X, Y\}$$

$$\mathcal{F}_H = \begin{cases} X \leftarrow f_X^H(u_X) & = u_X \\ Y \leftarrow f_Y^H(u_Y) & = u_Y \end{cases}$$

$$P(\mathbf{U}_H) = \begin{cases} P(U_X = 1) & = 0.5 \\ P(U_Y = 1) & = 0.2 \end{cases}$$

Indeed, this choice of \mathcal{M}_H is defined over X and Y . However, it does not seem like there is any connection between \mathcal{M}_H and \mathcal{M}_L from Example 3, even if $\mathbf{V}_H = \tau(\mathbf{V}_L)$. To verify this, we can compare the distributions induced by the two models. Note that while $P^{\mathcal{M}_L}(Y = 1 \mid A = 1, B = 1) \approx 0.853$, we see that $P^{\mathcal{M}_H}(Y = 1 \mid X = 1) = P^{\mathcal{M}_H}(Y = 1) = P^{\mathcal{M}_H}(U_Y = 1) = 0.2$. It seems that these two quantities should be related, as τ maps $(A = 1, B = 1)$ to $X = 1$, yet they are clearly not equal in the two models. This is similar for the causal effect $P^{\mathcal{M}_L}(Y_{A=1, B=1} = 1)$. Computing $P^{\mathcal{M}_H}(Y_{X=1} = 1) = P^{\mathcal{M}_H}(Y = 1) = 0.2$ actually yields the same result, which is clearly incorrect. In fact, it even seems that the causal relations are incorrect, as f_Y^H does not use X as an input. ■

Example 21 shows an example of \mathcal{M}_H that is a poor abstraction of \mathcal{M}_L despite the fact that it is defined over the space of \mathbf{V}_H mapped by τ . From the example, intuition tells us that a proper abstraction of \mathcal{M}_L should match \mathcal{M}_L in certain quantities, including observational, interventional, and counterfactual quantities. Specifically, there are quantities induced by \mathcal{M}_L that appear to have matching counterparts in \mathcal{M}_H based on τ . This notion is made concrete through the concept of Q - τ consistency (Def. 7).